

IBM Spectrum Control



Tables of Contents

Welcome	1
What's new	1
Change history	1
Discontinued features	9
Sponsor user program	10
Beta program	11
Collaborating with the team	11
Printable documentation	12
Product licenses	12
Definitions of licensing metrics	13
Names of equivalent licenses in previous releases	13
Actions that are available based on role	14
Licensing	15
How licensing values are calculated	18
Configuring license values for storage systems	19
Licensing examples	20
Getting started	21
Social media	22
Videos	22
Blogs	24
Product overview	26
Supported devices	27
Architecture	28
Interfaces for IBM Spectrum Control	29
Starting IBM Spectrum Control	29
Navigation	29
IBM Spectrum Storage Suite	32
IBM Virtual Storage Center	33
IBM Storage Insights	35
Replication products	36
Product updates and security fixes	37
Releases and downloads	37
Security	38
Subscribing to IBM announcements	39
End of support	39
Collaborating with the team	39
Reporting a problem	40
Key concepts	40
Data collection	41
Performance monitoring and troubleshooting	41
Applications, departments, and general groups	42
Storage optimization	43
Cloud configuration	43
Service classes	43
Capacity pools	45
Provisioning storage	46
Storage reclamation	47
Alerts and alert policies	48
Reporting	49
Rollup servers	49
Units of measurement for storage data	50
Role-based authorization	51
Fabrics and zones	51
Monitoring fabrics and zones in the GUI	52
Zones, zone aliases, and zone sets	52
Switch zoning capabilities	53
Zoning for Cisco MDS 9000 switches	53
Brocade switches in zones	53
Agents	53
SMI-S providers	53
Storage Resource agents	54
SNMP agents	54
REST APIs	54
Tutorials	55
Optimizing the performance of storage virtualizers	55
Tutorial: Analyzing and re-tiering volumes in pools on tier 1	56
Tutorial: Collocating volumes	58
Monitoring capacity usage at different levels of a business hierarchy	58
Tutorial: Comparing storage usage in each department	59
Using applications and subcomponents to monitor capacity and space usage	60

Tutorial: Viewing storage capacity and usage trends	61
Tutorial: Viewing NPIV connections between server ports and switch ports in a fabric	62
Tutorial: Exporting and uploading performance data for a SAN Volume Controller system	62
Tutorial: Comparing the performance of storage systems	64
Tutorial: Reviewing and updating your agentless servers	65
Tutorial: Troubleshooting performance	67
Tutorial: Monitoring IBM Spectrum Scale performance	68
Tutorial: Viewing the aggregated workload for an application	69
Tutorial: Identifying the source of slow drain problems caused by depletion of buffer credits	70
Tutorial: Identifying the locations of devices	72
Planning	72
Planning for installation	73
Planning to install IBM Spectrum Control in a Windows domain	75
Windows domain and local user accounts	76
Adding a computer to the Windows domain	76
Verifying that the Netlogon service is running	77
Installing Db2 by using a Windows domain user account	77
Creating a Windows domain common user account for IBM Spectrum Control	77
Granting Db2 SYSADM authority to a Windows domain user account	78
Required user privileges in installation scenarios	78
Planning for configuration	80
Planning for capacity	81
Planning for IBM Spectrum Control authentication and authorization	81
User names and passwords	82
User name and password requirements	82
Db2 user names and passwords	83
Required user roles for monitoring resources	83
Worksheet for user names and passwords	85
Ports used by IBM Spectrum Control	86
Planning for multipath subsystem device drivers	90
Planning for Storage Resource agents	90
Protocol support for Storage Resource agents	90
Planning for Internet Protocol Version 6	91
Planning to use LDAP	92
Planning for storage management	92
Planning for storage systems	93
Planning for monitoring storage systems	94
Planning for SMI-S providers	95
Planning for TagmaStore CIM agents	96
Planning for the native interface	97
Storage capacity of volumes	97
Planning for DS8000	98
Planning for SAN Volume Controller	98
Planning for IBM Spectrum Virtualize for Public Cloud	99
Planning for Storwize V7000	100
Planning for Storwize V7000 Unified	101
Planning for XIV storage systems	101
Planning for Dell EMC storage systems	102
Planning for Hitachi storage systems	103
Planning for NetApp device support	104
Planning for Pure Storage systems	106
Planning to monitor performance	108
Performance metrics	108
Performance configuration	109
Planning for switches and fabrics	110
Collecting data about fabrics	111
Planning for Brocade	111
Planning to use the BNA SMI agent	113
Support for Access Gateway switches and virtual fabrics	113
Planning for Cisco	113
Cisco SNMP agent	114
Configuration requirements for Cisco switches and directors	114
Migrating fabrics and switches	115
Planning for private switch networks	116
Planning for VMware	116
Configuring VMware in a IBM Spectrum Control environment	117
VMware capacity reports	117
Planning for files systems and volume managers	118
File systems	118
Networked file systems	118
Volume managers	118
Planning for PowerHASystemMirror for AIX	119
PowerHA SystemMirror for AIX environment	120
NAS support	120

Network-attached storage system requirements	120
Microsoft Cluster Server	121
Microsoft Cluster Server environment	121
Microsoft Cluster Server support	122
Planning for the Virtual I/O Server	122
Planning to monitor Db2	123
Sudo command privileges	124
Installing	124
IBM Spectrum Control components	124
Hardware requirements	125
Software requirements	125
Software requirements for operating systems	125
Software requirements for Storage Resource agents	125
Software requirements for the database repository	126
Web browser support	126
Software requirements for LDAP servers	126
Software requirements for CIM agents	126
Installing IBM Spectrum Control	126
Installation checklists for IBM Spectrum Control	127
Db2	128
Preparing to install Db2	129
Preparing to install Db2 on Windows	129
Preparing to install Db2 on UNIX or Linux	129
Installing Db2	130
Installing Db2 on Windows	130
Installing Db2 on AIX or Linux	131
Installing Db2 on AIX or Linux by using the command-line	132
Verifying that Db2 is installed correctly	133
Verifying Db2 installation using the command-line processor (CLP)	134
Verifying Db2 installation by using the First Steps tool	134
Licensing Db2	134
Starting the installation programs	135
Installing IBM Spectrum Control in a single-server environment	137
Installing IBM Spectrum Control in a single-server Windows environment	137
Installing IBM Spectrum Control in a single-server AIX or Linux environment	137
Installing IBM Spectrum Control in a single-server environment by using silent mode	138
Editing the response file	140
Installing IBM Spectrum Control in a multiple-server environment	142
Installing IBM Spectrum Control with a remote database by using the installation program	143
Installing IBM Spectrum Control with a remote database by using silent mode	144
Accessing IBM Spectrum Control server from a remote computer using Command Line Interface	147
Installing IBM Spectrum Control on a Windows domain	147
Verifying the connection to the domain controller computer by using the Dcdiag tool	149
Installing IBM Spectrum Control and associated products using minimal space on the Windows C: drive	150
Verifying the installation	150
Reviewing the log files to resolve installation issues	152
Changing languages	152
Changing the language of the IBM Spectrum Control GUI	152
Changing the operating system language for Windows	153
Changing the operating system language for AIX	153
Adding an installation license	153
Adding an installation license using the installation program	154
Adding an installation license using silent mode	154
Installing Storage Resource agents	154
Installing Storage Resource agents by using a command	155
Installing IBM Cognos Analytics	156
Installing Cognos Analytics in a single Windows environment	156
Installing Cognos Analytics in a single Linux environment	157
Installing Cognos Analytics in a single AIX environment	158
Configuring IBM Cognos Analytics	159
Copying JDBC files for Db2 on Windows	159
Copying Db2 files and linking the Db2 library on Linux and AIX	160
Creating a content store and starting Cognos Analytics on Windows	160
Creating a content store and starting Cognos Analytics on Linux or AIX	161
Creating Cognos Analytics data source for the IBM Spectrum Control database	162
Using Cognos Analytics with multiple remote IBM Spectrum Control databases	163
Importing IBM Spectrum Control pre-defined reports package	164
Configuring access to the Cognos Analytics reporting tool	165
Reinstalling the software if a failure occurs	165
Taking the first steps after installation	166
Uninstalling	166
Uninstalling IBM Spectrum Control in a single-server environment	167
Uninstalling IBM Spectrum Control in a multiple-server environment	167
Uninstalling IBM Spectrum Control by using silent mode	168
Manually uninstalling IBM Spectrum Control components on AIX or Linux	169

Manually uninstalling IBM Spectrum Control components on Windows	169
Uninstalling Storage Resource agents	170
Uninstalling Storage Resource agents manually	171
Deleting Storage Resource agent registry entries after a failed installation or uninstallation	171
Removing servers	172
Uninstalling Db2	172
Uninstalling Db2 on AIX or Linux	172
Uninstalling IBM Cognos Analytics	174
Upgrading and migrating	174
Preparing to upgrade	175
Backups	177
Example of upgrading Db2 in a Windows environment	177
Example of upgrading Db2 in a AIX or Linux environment	178
Upgrading IBM Spectrum Control in a single-server environment	179
Upgrading IBM Spectrum Control in a single-server Windows environment by using the installation program	180
Upgrading IBM Spectrum Control in a single-server AIX or Linux environment by using the installation program	180
Upgrading IBM Spectrum Control in a single-server environment by using silent mode	181
Editing the upgrade response file	182
Upgrading IBM Spectrum Control in a multiple-server environment	183
Upgrading to IBM Spectrum Control 5.3.0 or later with a remote database by using the installation program	183
Upgrading IBM Spectrum Control with a remote database by using silent mode	183
Upgrading Storage Resource agents	185
Starting agent upgrades	186
Scheduling agent upgrades	186
Upgrading Storage Resource agents by using a command	186
Upgrading SMI-S providers for storage systems	187
Upgrading CIM agents	187
Migrating alert definitions to alert policies	187
Scenario: Changing from LDAP and Local OS authentication to only LDAP or Local OS authentication after an upgrade	188
Configuring	188
Starting IBM Spectrum Control	189
Configuring history and data retention	190
Configuring user authentication	191
Authorizing users	191
Role-based authorization	192
Assigning a role to a group	192
Determining the groups to which a user belongs	193
Modifying the authentication mechanism	193
Actions that are available based on role	194
Managing authentication	195
Changing from operating system to LDAP authentication	196
Changing from LDAP to operating system authentication	198
Configuring user authentication alternatives	198
Enabling secure communication between IBM Spectrum Control and the LDAP repository	198
Disabling secure communication between IBM Spectrum Control and the LDAP repository	199
Exporting SSL certificate from the IBM Security Directory Server to a file	199
Exporting SSL certificate from the Microsoft Active Directory to a file	200
Adding the SSL certificate for the LDAP server to the web server keystore using the IKEYCMD command	200
Using the ldapEntityType element for advanced LDAP configuration	201
Adding customized text to the logon page	202
Managing a SAN without agents	202
Setting timeout values for the Device server	203
Configuring Service Location Protocol	204
Router configuration	204
SLP directory agent configuration	204
Environment configuration	204
SLP registration and the slptool command	205
SLP discovery	205
Configuring IP addressing	205
Configuring IBM Spectrum Control with multiple IP addresses	205
Changing the HOSTS file	206
Deploying Storage Resource agents	207
Deployment guidelines and limitations for Storage Resource agents	207
Creating a certificate for SSH protocol	210
Replacing default SSL certificates for the Data server and Storage Resource agents with custom SSL certificates	213
Configuration guidelines for 500 or more agents	215
Including a Storage Resource agent with a server golden image	216
Checking for a fully qualified host name	216
Checking for a fully qualified host name for AIX systems	216
Checking for a fully qualified host name for Linux systems	217
Checking for a fully qualified host name for Oracle Solaris	217
Checking for a fully qualified host name for Windows systems	217
Granting local administrative privileges to a domain account	218
Importing authentication information for a Storage Resource agent	218
Installing and configuring the IBM Spectrum Control server with multiple NIC cards	218
Replacing the default SSL certificate for the Device, Alert, or Web server with a self-signed certificate	219

Replacing the default SSL certificate for the Device, Alert, or Web server with a certificate from an external certificate authority	220
Updating IBM Spectrum Control data collector trusted certificates after replacing default SSL certificate for the Device server with a self-signed certificate OR an external certificate	221
Replacing the default SSL certificate for the Export server	222
Generating a new default self-signed SSL certificate for the Export server	222
Enabling TLS 1.0 and 1.1 for ports	223
Enabling TLS 1.1 and 1.0 for IBM Spectrum Control ports	223
Configuring Db2, AIX, and Linux for IPv6-only environment	224
Administering	225
Administering resources and data sources	225
Storage systems	226
Viewing information about storage systems	226
Updating the credentials for storage systems	226
Updating the credentials for a System Storage DS8000 storage system	227
Updating the credentials for an XIV or IBM Spectrum Accelerate	227
Updating the credentials for storage systems that run IBM Spectrum Virtualize	227
Updating the credentials for a Storwize V7000 Unified storage system	228
Updating the credentials for a storage system that is managed by a CIM agent	229
Testing the connection to a storage system	229
Collecting CIM agent logs	229
Removing storage systems	230
Hypervisors and VMware data sources	230
Checking permissions to browse data stores	230
Viewing information about hypervisors	230
Updating the credentials for a hypervisor	231
Removing hypervisors and VMware data sources	231
Switches and fabrics	231
Viewing information about switches and fabrics	231
Updating the connection information for switches and fabrics	232
Updating the connection information for a switch	232
Updating the connection information for a fabric	233
Testing the connection to a switch or fabric	233
Removing switches and fabrics	234
Servers and Storage Resource agents	234
Fixing deployments	235
Canceling deployments	236
Modifying deployment schedules	236
Viewing information about Storage Resource agents	236
Viewing Storage Resource agent log files	237
Disabling Storage Resource agents	237
Enabling Storage Resource agents	238
Testing the connection with a Storage Resource agent	238
Changing credentials for Storage Resource agents	238
Collecting service data	239
Enabling or disabling scripts for Storage Resource agents	239
Enabling or disabling the monitoring of fabrics by Storage Resource agents	240
Using the help command for Storage Resource agents	240
Removing servers	240
Registering a Storage Resource agent with a different IBM Spectrum Control server	241
Manually changing the Windows service logon	241
Deployment guidelines and limitations for Storage Resource agents	241
SMI-S providers	244
Verifying that an SMI agent is running	244
Replacing an SMI agent for block storage systems, fabrics, and switches	245
Interop namespaces for SMI-S providers for switches and storage systems	245
SNMP agents	246
Displaying information about an SNMP agent	246
Removing an SNMP agent	246
Starting and stopping the IBM Spectrum Control servers	247
Starting the servers by using the GUI	247
Starting the servers by using scripts	247
Stopping the servers by using the GUI	248
Stopping the servers by using scripts	249
Checking the version and license of IBM Spectrum Control	249
Checking IBM Spectrum Control status	250
Troubleshooting problems with the IBM Spectrum Control component and servers	250
Packaging and sending log files from the System Management page	250
Increasing the memory allocation for the Data server	251
Increasing the memory allocation for the Data server that is running on AIX	252
Increasing the memory allocation for the Data server that is running on Linux	252
Increasing memory allocation for Data server that is running on Windows	253
Changing passwords	253
Changing passwords by using the password tool	253
Single server installation where components use the same logon credentials	254

Single-server installation where components use different logon credentials	255
Multiple-server installation where Db2 is remote	256
Changing passwords on AIX and Linux systems using the Command Line Interface (CLI)	257
Changing passwords on Windows systems from the Command Line Interface (CLI)	258
Granting local administrative privileges to a domain account	258
Collecting diagnostic information about IBM Spectrum Control	259
Service tool overview	259
Packaging log files from the command line and sending them to IBM Support	260
Creating a compressed file for a Storage Resource agent	261
How to customize the service tool	261
Administering the IBM Spectrum Control database	263
Backing up the database	263
Comparison of database backup methods	264
Backing up the database offline using the command line	265
Backing up the database offline using IBM® Data Studio client	265
Backing up the database online using IBM® Data Studio client	266
Restoring the database	267
Restoring the database using the command line	267
Restoring the database using IBM Data Studio client	267
Disaster recovery	268
Maintaining and improving the performance of the database	268
Updating database statistics	268
Reorganizing database tables	269
Customizing the reorganization function of the database maintenance tool	269
Repository copy tool	270
Exporting repository data	270
Administering Db2	271
Using the command line on UNIX and Linux	271
Manually starting Db2 on Windows	272
Manually stopping Db2 on Windows	272
Starting the IBM Data Studio full client	272
Monitoring Db2	273
Managing resources	273
Resources that you can monitor	274
Adding resources	277
Required user roles for monitoring resources	277
Adding storage systems	279
Hitachi	281
Installing the Hitachi Export Tool	282
Adding and configuring NetApp resources	283
Configuring IBM Spectrum Virtualize for Public Cloud for monitoring	284
Monitoring IBM Spectrum Virtualize for Public Cloud with on-premises data collection (Site to Site VPN IPsec)	284
User roles for collecting performance metadata from IBM Spectrum Virtualize	285
Monitoring IBM Spectrum Scale without requiring root privileges	285
Verifying that probe data can be collected for object storage	286
Configuring the collection of performance data for IBM Spectrum Scale	286
Configuring OpenStack access to monitor the object storage system	287
Adding fabrics and switches	288
Data sources for switches and fabrics	289
Configuring Brocade switches for monitoring	289
Adding hypervisors	290
Adding ESX and ESXi hypervisors	290
Checking permissions to browse data stores	291
Adding vCenter Server systems	291
Adding servers	292
Agentless servers	292
Adding servers with Storage Resource agents	292
File List	293
Adding rollup servers	293
Restrictions for rollup servers	294
Removing resources	294
Collecting data	295
Collecting asset and status data	295
Creating probes	296
How automated probes are scheduled	296
Verifying that a probe is running for a resource	297
Modifying probes	297
Configuring alerts for probes	298
Starting probes	298
Viewing probes for a specific resource	299
Viewing probe logs	299
Collecting performance data	299
Creating performance monitors	300
Verifying that a performance monitor is running for a resource	301

Modifying performance monitors	301
Configuring alerts for performance monitors	302
Starting and stopping performance monitors	302
Viewing performance monitors	303
Viewing all performance monitors	303
Viewing performance monitors for specific resources	304
Viewing performance monitors logs	304
Collecting information about shares on storage systems	305
Collecting information about the sizes of snapshots in IBM Spectrum Scale	305
Alerting	305
How alerts work	306
Viewing and administering alerts	307
Viewing and administering alert definitions	308
Alert policies	308
Defining notification settings for alerts	309
Alert severities	310
Defining alerts	310
Defining alerts for resources	311
Defining alert definitions for general attributes and capacity changes	311
Defining alert definitions for performance changes	312
Defining custom alerts for resources	314
Scenarios for custom alerts	315
Defining alerts for applications	319
Defining application alerts for attribute and capacity changes	320
Defining application alerts for performance metrics	320
Defining custom alerts for applications	322
Defining alerts for general groups	323
Defining general group alerts for attribute and capacity changes	323
Defining general group alerts for performance metrics	324
Defining custom alerts for general groups	325
Configuring alert notifications	326
Configuring email alert notifications	327
Configuring SNMP alert notifications	327
Configuring Tivoli Netcool/OMNIBus alert notifications	328
Triggering conditions for alerts	328
Triggering conditions for storage system alerts	329
Triggering conditions for storage system internal resource alerts	332
Triggering conditions for hypervisor alerts	340
Triggering conditions for switch alerts	343
Triggering conditions for fabric alerts	344
Triggering conditions for server alerts	345
Alert notifications and actions	349
How scripts are run	350
Viewing information about resources	351
How information is organized	352
Dashboard view	353
Viewing information about top-level resources	355
Viewing information about internal, object, and related resources	355
Overview charts	356
Capacity by pool	358
Capacity by volume	359
Viewing information about enclosures	360
Monitoring the status and condition of resources	361
How the condition of a resource is determined	362
Viewing the overall condition of resources	364
Viewing the condition of specific types of resources	364
Viewing the status of resources	365
Acknowledging the condition and status of resources	365
Acknowledging the condition of top-level resources	365
Acknowledging the status of internal resources	366
Monitoring vaults in IBM Cloud Object Storage	367
Monitoring access to vaults in IBM Cloud Object Storage	367
Investigating vaults that cannot be accessed	367
Monitoring vaults that are at risk of access failure	368
Calculating the failure tolerance for vaults	369
Monitoring the performance of resources	369
Viewing performance information	370
How performance information is displayed	372
Controls	373
Resources in the chart legend	374
Saving URLs for performance views	375
Identifying performance issues for IBM Spectrum Virtualize storage systems	376
IBM Spectrum Virtualize guideline values for key performance indicators	376
Viewing performance alerts	378
Exporting performance data for storage systems and fabrics	378

Exporting performance data by using the GUI	379
Exporting performance data by using a script	379
Performance metrics	380
Performance metrics for DS8000	380
Performance metrics for resources that run IBM Spectrum Virtualize	385
Performance metrics for XIV, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R	395
Performance metrics for IBM Spectrum Scale	397
Performance metrics for IBM FlashSystem 900	398
Performance metrics for Dell EMC storage systems	400
Performance metrics for Hitachi VSP storage systems	404
Performance metrics for NetApp storage systems	405
Performance metrics for Pure storage systems	407
Performance metrics for other storage systems	408
Performance metrics for switches	409
Monitoring the capacity of resources	411
Setting capacity limits	412
Tutorial: Investigating compliance with the capacity limit	413
Removing the capacity limit	414
Capacity limit metrics	414
Viewing capacity information	415
How capacity information is displayed	416
Controls for capacity views	417
Resources in the capacity chart legend	418
Investigating capacity trends for block storage systems	418
Investigating capacity trends for block storage pools	419
Investigating capacity trends for volumes	420
Investigating capacity trends for file systems	420
Investigating capacity trends for file system pools	421
Investigating capacity trends for filesets	422
Investigating capacity trends for containers	423
Investigating capacity trends for tiers	424
Investigating capacity trends for servers	424
Creating bookmarks for URLs of capacity views	424
Viewing the capacity of external storage	425
Identifying shortfall before data is recalled from external storage	426
Viewing capacity alerts and violations	426
Capacity metrics	426
Capacity metrics for block storage systems	427
Key capacity concepts	435
Key storage values for pools	441
Capacity metrics for file storage systems	441
Capacity metrics for object storage systems	442
Capacity metrics for tiers	443
Monitoring and administering applications, departments, and general groups	443
Applications	446
Creating an application hierarchy	447
Creating applications	447
Using filters to assign resources to applications	448
Adding subcomponents	449
Adding resources manually to applications	450
Adding resources with the command line interface	450
Viewing information about applications	450
Removing applications and subcomponents	451
Removing resources from applications	451
Departments	452
Creating a department hierarchy	453
Creating departments	453
Adding and creating subdepartments	454
Adding applications to departments	455
Viewing information about departments	455
Removing departments and subdepartments	455
Removing applications from a department	456
Application and department hierarchies	456
Exporting information about applications and departments	457
General groups	458
Creating a general group hierarchy	458
Creating general groups and adding resources	458
Adding general groups as subgroups	459
Adding resources to general groups	459
Viewing and administering general groups	460
Provisioning storage	462
Configuring IBM Spectrum Control for provisioning	462
Configuring service classes	463

Creating service classes	463
Viewing the details of service classes	464
Modifying service classes	464
Deleting service classes	465
Tagging resources to satisfy custom requirements	465
Configuring capacity pools	466
Adding resources to capacity pools	466
Viewing the details of capacity pools	467
Modifying capacity pools	468
Deleting capacity pools	468
Block storage: Calculating available capacity and determining the placement of volumes	468
File storage: Calculating available capacity and determining the placement of shares	470
Changing the default host definition for provisioned storage	471
Provisioning storage with the IBM Spectrum Control GUI	472
Provisioning volumes to servers	472
Provisioning volumes to hypervisors	473
Provisioning volumes to a server cluster	473
Provisioning volumes to a hypervisor cluster	473
Provisioning shares to servers	474
Provisioning shares to hypervisors	474
Provisioning shares to a server cluster	474
Provisioning shares to a hypervisor cluster	475
Provisioning with zone control	475
Optimizing storage tiering	475
Investigating the capacity of tiered storage	477
Setting the tier level of storage pools	478
Renaming tiers	478
Analyzing tiering by servers	478
Analyzing tiering by hypervisors	479
Analyzing tiering by storage virtualizers	479
Analyzing tiering by storage pools	479
Analyzing tiering by volumes	479
Tiering volumes by I/O density and I/O rate	479
Modifying the criteria for analyzing tiering	480
Optimizing storage pools	480
Balancing pools	481
Modifying the criteria for balancing pools	481
Criteria for identifying the pools that require balancing	482
Reclaiming storage	482
Viewing storage reclamation	483
Viewing storage reclamation by storage system	483
Excluding volumes from reclamation analysis	484
Transforming and migrating volumes	484
Modifying the period for analyzing performance data	485
Modifying the properties of resources	485
Properties of resources that can be modified	486
Opening the management GUI for a resource	487
Exporting information to a file	487
Customizing lists	488
Filtering lists	488
Sorting lists	489
Showing, hiding, and reordering columns in lists	489
Managing tasks	489
Viewing tasks	490
Viewing all tasks	490
Viewing pending tasks	491
Viewing tasks for a type of resource	491
Viewing tasks for a specific resource	491
Viewing task details	492
Viewing task logs	492
Setting the number of task runs that are displayed	493
Managing tasks for provisioning	493
Creating provisioning tasks	493
Running provisioning tasks	494
Exporting provisioning task information	494
Managing tasks for tiering storage, balancing pools, and transforming storage	494
Creating analysis tasks	495
Running analysis tasks	496
Implementing recommendations to optimize storage	496
Implementing recommendations immediately	496
Scheduling the implementation of recommendations	497
Pausing, resuming, and canceling the implementation of recommendations	498
Renaming tasks	498
Deleting tasks	499
Planning copy data resources	499

Monitoring copy data resources	499
REST APIs	500
Retrieve data by using a REST API command line utility	500
Retrieve data by using REST APIs with a web browser	502
Reporting	503
Creating predefined capacity reports	505
Tutorial: Creating a predefined capacity report about storage systems	508
Tutorial: Creating a predefined capacity report about pools	509
Tutorial: Creating a predefined capacity report about tiered pools	510
Tutorial: Creating a predefined capacity report about the volumes assigned to servers	510
Tutorial: Creating a predefined capacity report about managed disks by storage systems	511
Adding resources to applications to generate large reports	511
Creating predefined inventory reports	512
Tutorial: Creating an inventory report about block storage systems	513
Creating custom reports	513
Tutorials: Creating custom capacity and performance reports for applications	515
Creating the application for the capacity report	515
Creating the capacity report for the application	516
Creating the performance report for the application	517
Custom capacity and performance view reports	517
Creating chargeback and consumer reports	519
Creating chargeback reports	520
Chargeback reports	521
Creating consumer reports	522
Consumer reports	523
Creating summary reports of the storage capacity	523
Running reports	524
Editing reports	524
Deleting reports	524
Configuring the email server	524
Investigating issues with reports	525
Types of predefined capacity and inventory reports	526
Reports FAQ	526
Capacity metrics for reports	527
Capacity metrics for chargeback and consumer reports	527
Using the REST API to generate reports	529
Retrieve data by using a REST API command line utility	529
Retrieve data by using REST APIs with a web browser	530
Exporting information about resources	532
Exporting information about capacity shortages in pools	532
Exporting information about the input/output performance of volumes	533
Reporting with Cognos Analytics	533
Getting started	534
Predefined reports	534
Predefined reports listed by package	535
Custom reports	538
Viewing predefined reports	538
Creating custom reports	539
Viewing and creating reports in the Cognos Analytics reporting tool	539
Searching in your report data	539
Scenarios	540
Analyzing the volumes on the most active hypervisors	540
Monitoring daily the performance of volumes	541
Investigating a degradation in the performance of a storage pool	541
Identifying the relationships between a storage system and other resources	542
Configuring	542
Setting options for reports	542
Scheduling a report to run with predefined options	543
Creating a copy of a predefined report	543
Starting the Cognos Configuration GUI	544
Sending reports by email	544
Specifying locations for saving reports	544
Specifying a root directory for saving reports	545
Specifying file system locations for reports	545
Specifying a report to save the report output to a file system	546
Setting up multiple connections to the database repository	546
Customizing the logo and title of reports	547
Creating a folder in the Cognos Analytics reporting tool	547
Creating a backup copy of a report	547
Changing the logo for all reports	547
Modifying the title of a report	548
Predefined reports about resource relationships	548
Running the Storage Resource Relationships Summary report	548
Storage Resource Relationships Summary report	549
Running the Storage Resource Relationships Summary (Configurable) report	549

Storage Resource Relationships Summary (Configurable) report	550
Running the File System to Volume Relationships report	551
File System to Volume Relationships report	551
Predefined reports about switches and switch ports	552
Predefined reports about switches	552
Running the Availability of Switch Ports report	553
Availability of Switch Ports report	553
Running the Performance of One Switch report	554
Performance of One Switch report	554
Running the Compare Performance of Multiple Switches report	555
Compare Performance of Multiple Switches report	555
Running the Compare Performance of One Switch over Time Ranges report	556
Compare Performance of One Switch over Time Ranges report	556
Predefined reports about switch ports	556
Running the Most Active Switch Ports report	557
Most Active Switch Ports report	557
Running the Performance of One Switch Port report	558
Performance of One Switch Port report	558
Running the Compare Performance of Multiple Switch Ports report	559
Compare Performance of Multiple Switch Ports report	560
Running the Compare Performance of One Switch Port over Time Ranges report	560
Compare Performance of One Switch Port over Time Ranges report	561
Predefined reports about groups	561
Running the Capacity of One Group report	561
Capacity of One Group report	562
Running the Groups Capacity report	563
Groups Capacity report	563
Predefined reports about hypervisors	564
Running the Hypervisors Capacity report	564
Hypervisors Capacity report	565
Running the Most Active Hypervisors report	565
Most Active Hypervisors report	566
Running the Summarized Performance of Volumes by Hypervisor report	566
Summarized Performance of Volumes by Hypervisor report	567
Running the Performance of Volumes by Hypervisor report	567
Performance of Volumes by Hypervisor report	568
Running the Hypervisor Data Stores Capacity report	568
Hypervisor Data Stores Capacity report	569
Running the Hypervisor Disks Capacity report	569
Hypervisor Disks Capacity report	570
Predefined reports about servers and file systems on servers	570
Predefined reports about servers	571
Running the Servers Capacity report	571
Servers Capacity report	571
Running the Most Active Servers report	572
Most Active Servers report	572
Running the Summarized Performance of Volumes by Server report	573
Summarized Performance of Volumes by Server report	573
Running the Performance of Volumes by Server report	574
Performance of Volumes by Server report	575
Running the Server Disks Capacity report	575
Server Disks Capacity report	576
Predefined reports about file systems on servers	576
Running the File Systems Capacity report	576
File Systems Capacity report	577
Predefined reports about storage systems and components	577
Predefined reports about storage systems	578
Running the Storage Systems Capacity report	578
Storage Systems Capacity report	578
Running the Storage Systems Historical Capacity report	580
Storage Systems Historical Capacity report	580
Running the Most Active Storage Systems report	581
Most Active Storage Systems report	582
Running the Performance of One Storage System report	583
Performance of One Storage System report	583
Running the Compare Performance of Multiple Storage Systems report	584
Compare Performance of Multiple Storage Systems report	584
Running the Compare Performance of One Storage System over Time Ranges report	585
Compare Performance of One Storage System over Time Ranges report	585
Running the Performance Data Export report	586
Performance Data Export report	586
Predefined reports about controllers, modules, and nodes	586

Running the Most Active Controllers or Modules report	587
Most Active Controllers or Modules report	587
Running the Performance of One Controller or Module report	588
Performance of One Controller or Module report	588
Running the Most Active Nodes report	589
Most Active Nodes report	589
Running the Performance of One Node report	590
Performance of One Node report	591
Running the Compare Performance of Multiple Controllers and Modules report	592
Compare Performance of Multiple Controllers and Modules report	592
Running the Compare Performance of Multiple Nodes report	592
Compare Performance of Multiple Nodes report	593
Predefined reports about disks	593
Running the Disks Capacity report	594
Disks Capacity report	594
Running the Most Active Disks report	595
Most Active Disks report	595
Running the Performance of One Disk report	596
Performance of One Disk report	596
Running the Compare Performance of Multiple Disks report	597
Compare Performance of Multiple Disks report	598
Predefined reports about host connections	598
Running the Most Active Host Connections report	598
Most Active Host Connections report	599
Running the Summarized Performance of Volumes by Host Connection report	600
Summarized Performance of Volumes by Host Connection report	600
Running the Performance of Volumes by Host Connection report	601
Performance of Volumes by Host Connection report	601
Predefined reports about I/O groups	602
Running the Most Active IO Groups report	602
Most Active IO Groups report	603
Running the Performance of One IO Group report	603
Performance of One IO Group report	604
Running the Compare Performance of Multiple IO Groups report	605
Compare Performance of Multiple IO Groups report	605
Predefined reports about managed disks	606
Running the Managed Disks Capacity report	606
Managed Disks Capacity report	606
Running the Most Active Managed Disks report	607
Most Active Managed Disks report	608
Running the Performance of One Managed Disk	608
Performance of One Managed Disk report	609
Running the Compare Performance of Multiple Managed Disks report	610
Compare Performance of Multiple Managed Disks report	610
Predefined reports about storage pools	611
Running the Pools Capacity report	611
Pools Capacity report	611
Running the Pools Historical Capacity report	613
Pools Historical Capacity report	614
Running the Most Active Pools report	615
Most Active Pools report	615
Running the Performance of One Pool report	616
Performance of One Pool report	616
Running the Compare Performance of Multiple Pools	617
Compare Performance of Multiple Pools report	618
Running the Compare Performance of One Pool over Time Ranges report	618
Compare Performance of One Pool over Time Ranges report	618
Predefined reports about ports	619
Running the Most Active Ports report	619
Most Active Ports report	620
Running the Performance of One Port report	620
Performance of One Port report	621
Running the Compare Performance of Multiple Ports report	622
Compare Performance of Multiple Ports report	622
Predefined reports about RAID arrays	623
Running the Most Active RAID Arrays report	623
Most Active RAID Arrays report	623
Running the Performance of One RAID Array report	624
Performance of One RAID Array report	625
Running the Compare Performance of Multiple RAID Arrays report	625
Compare Performance of Multiple RAID Arrays report	626
Predefined reports about volumes	626

Running the Volumes Capacity report	627
Volumes Capacity report	627
Running the Volumes Historical Capacity report	629
Volumes Historical Capacity report	629
Running the Most Active Volumes report	630
Most Active Volumes report	631
Running the Performance of One Volume report	631
Performance of One Volume report	632
Running the Compare Performance of Multiple Volumes report	633
Compare Performance of Multiple Volumes report	633
Predefined reports about storage tiering	634
Running the Storage Resource Group - VDisk Workload Activity Details report	635
Storage Resource Group - VDisk Workload Activity Details report	635
Running the VDisk Details report	638
VDisk Details report	638
Running the MDisk Groups - VDisk Workload Activity report	639
MDisk Groups - VDisk Workload Activity report	640
Running the Storage Virtualizer - VDisk Workload Activity report	642
Storage Virtualizer - VDisk Workload Activity report	642
Running the Storage Resource Group - VDisk Workload Activity report	645
Storage Resource Group - VDisk Workload Activity report	645
Running the MDisk Group - VDisk Workload Activity Details report	647
MDisk Group - VDisk Workload Activity Details report	648
Running the MDisk Group Details report	650
MDisk Group Details report	650
Running the MDisk Groups - Workload Activity report	652
MDisk Groups - Workload Activity report	652
Read I/O capability formula	655
Changing threshold values in storage tier reports	656
Threshold values in storage tier reports	656
Custom reports about performance	657
Creating custom performance reports	657
Data and properties in performance reports	658
Data for storage systems in performance reports	658
Performance metrics for SAN Volume Controller and Storwize systems	660
Performance metrics for DS8000 storage systems	664
Performance metrics for XIV systems and IBM Spectrum Accelerate	668
Performance metrics for block server systems	669
Data for storage system volumes in performance reports	670
Performance metrics for volumes on SAN Volume Controller and Storwize systems	673
Performance metrics for volumes on DS8000 storage systems	675
Performance metrics for storage system volumes on XIV systems and IBM Spectrum Accelerate	676
Performance metrics for storage system volumes on block servers	678
Data for storage system ports in performance reports	679
Performance metrics for ports on SAN Volume Controller and Storwize systems	679
Performance metrics for ports on DS8000 and DS6000 storage systems	681
Performance metrics for ports on an XIV	683
Performance metrics for storage system ports on block servers	684
Data for storage system RAID arrays in performance reports	684
Performance metrics for storage system arrays	685
Data for storage system controllers in performance reports	687
Performance metrics for storage system controllers on DS8000 storage systems	687
Performance metrics for storage system controllers on block servers	689
Data for storage system modules in performance reports	690
Performance metrics for storage system modules on XIV systems and IBM Spectrum Accelerate	690
Data for storage system pools in performance reports	692
Performance metrics for pools on SAN Volume Controller and Storwize systems	694
Performance metrics for pools on DS8000 storage systems	696
Performance metrics for pools on XIV systems and IBM Spectrum Accelerate	698
Data for managed disks on storage systems in performance reports	699
Performance metrics for managed disks on SAN Volume Controller and Storwize systems	700
Data for storage system nodes in performance reports	701
Performance metrics for nodes on SAN Volume Controller and Storwize systems	701
Data for storage system I/O groups in performance reports	705
Performance metrics for I/O groups on SAN Volume Controller and Storwize systems	706
Data for local disks on storage systems in performance reports	710
Performance metrics for local disks on storage systems	711
Data for storage host connections in performance reports	712
Performance metrics for host connections on SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified systems	712
Performance metrics for host connections on DS8000 storage systems	714
Performance metrics for storage host connections on an XIV	715
Data for switches in performance reports	716

Performance metrics for switches	717
Data for switch ports in performance reports	719
Performance metrics for switch ports	719
Custom reports about capacity and relationships	721
Creating custom capacity and relationship reports	721
Adding filters to reports	722
Filters for servers in capacity and relationship reports	723
Filters for hypervisors in capacity and relationship reports	723
Filters for storage virtualizers in capacity and relationship reports	724
Filters for storage systems in capacity and relationship reports	725
Filters for storage resource groups in capacity and relationship reports	725
Data and properties in capacity and relationship reports	725
Data for clusters in capacity and relationship reports	726
Data for servers in capacity and relationship reports	727
Data for file systems on servers in capacity and relationship reports	728
Data for logical volumes on servers in capacity and relationship reports	729
Data for volume groups on servers in capacity and relationship reports	729
Data for server disks in capacity and relationship reports	730
Data for server groups in capacity and relationship reports	731
Data for server controllers in capacity and relationship reports	731
Data for multipath drivers on servers in capacity and relationship reports	731
Data for file system groups on servers in capacity and relationship reports	732
Data for hypervisors in capacity and relationship reports	732
Data for clusters on hypervisors in capacity and relationship reports	733
Data for file systems on hypervisors in capacity and relationship reports	733
Data for data stores on hypervisors in capacity and relationship reports	734
Data for server disks on hypervisors in capacity and relationship reports	735
Data for hypervisor controllers in capacity and relationship reports	735
Data for multipath drivers on hypervisors in capacity and relationship reports	736
Data for virtual machines managed by hypervisors in capacity and relationship reports	736
Data for network-attached storage systems in capacity and relationship reports	737
Data for file systems on network-attached storage systems in capacity and relationship reports	738
Data for network-attached storage exports in capacity and relationship reports	738
Data for logical volumes on network-attached storage systems in capacity and relationship reports	739
Data for network-attached storage disks in capacity and relationship reports	739
Data for pools on Storwize V7000 Unified systems in capacity and relationship reports	740
Data for Storwize V7000 Unified filesets in capacity and relationship reports	741
Data for storage virtualizers in capacity and relationship reports	741
Data for storage virtualizer pools in capacity and relationship reports	743
Data for storage virtualizer volumes in capacity and relationship reports	746
Data for managed disks on storage virtualizers in capacity and relationship reports	749
Data for local disks on storage virtualizers in capacity and relationship reports	750
Data for storage virtualizer groups in capacity and relationship reports	751
Data for storage systems in capacity and relationship reports	751
Data for storage system pools in capacity and relationship reports	753
Data for storage system volumes in capacity and relationship reports	756
Data for storage system disks in capacity and relationship reports	760
Data for storage system groups in capacity and relationship reports	761
Data for managed disks on storage systems in capacity and relationship reports	761
Data for storage resource groups in capacity and relationship reports	762
Data for groups in capacity and relationship reports	762
Data for switches in capacity and relationship reports	762
Data for switch ports in capacity and relationship reports	764
Data for fabrics in capacity and relationship reports	764
Troubleshooting Cognos Analytics reports	765
Showing the package version number and build ID	767
Reference	767
Return codes used by Storage Resource agent	767
Agent types for monitoring fabrics and switches	769
Supported storage systems providing full disk encryption and solid-state drives	769
Performance counters	770
Command-line interface	788
CLI requirements	789
Command modes	789
Conventions used in this guide	789
Syntax diagram conventions	790
tpctool command	790
actzs	796
addza	797
addzports	797
addzone	798
addzoneports	798

assignvol	799
autosetarray	800
catdscfg	802
chexport	803
chfs	805
chfset	807
chkquota	808
chwcache	809
chwcachesource	810
ckzone	811
ckzs	812
commit	812
ctlwcache	813
deactzs	815
encrypt	815
getdscfg	815
getrpt	816
linkfset	817
lsappgroup	818
lsappgroupmembers	819
lsarray	820
lsbackenddisktypes	822
lsbackendraidthypes	823
lsbackendtypes	823
lscluster	824
lscomp	825
lscounters	826
lsdeptgroup	827
lsdeptgroupmembers	829
lsdev	830
lsdevp	831
lsdisk	832
lsexport	833
lsextent	834
lsfcpath	835
lsfs	835
lsfset	837
lshtype	839
lsoptschedules	840
lsmetrics	840
lsnode	841
lsnsd	843
lspool	844
lsport	845
lsquota	846
lssrg	847
lssrgmembers	848
lssvr	849
lssvrdisk	850
lsswitch	850
lstime	851
lstype	852
lsvm	853
lsvmdisk	854
lsvol	854
lsvolgroup	855
lszone	856
lszs	857
mkappgroup	858
mkdeptgroup	859
mkexport	860
mkfs	862
mkfset	864
mksrg	865
mkzone	866
mkzs	866
modifyappgroup	867
modifyappgroupviafile	870
modifydeptgroup	872
modifydeptgroupviafile	874
modifyfsg	876
mountfs	877

rmappgroup	878
rmbackenddisktype	879
rmbackendraidtype	880
rmbackendtype	881
rmdeptgroup	882
rmexport	883
rmfs	883
rmfset	884
rmsrg	885
rmza	886
rmzaports	886
rmzone	887
rmzoneports	888
rmzs	888
rollback	889
runoptschedule	889
setarray	890
setbackenddisktype	891
setbackendraidtype	892
setbackendtype	893
setdscfg	895
setquota	895
showoptresults	897
start	898
unassignvol	898
unlinkfset	899
unmountfs	900
updatesrg	901
Command aliases	902
Parameter aliases	902
agent.sh command	903
dataCollector command	903
Configuration files	903
server.config file	904
scheduler.config file	905
TPCD.config file	906
Specifying the tablespace size for IBM Spectrum Control	906
agent.config file	907
Log files	907
Default locations of log files	907
Script parameters	908
Opening IBM Spectrum Control on Windows operating systems	910
Opening IBM Spectrum Control GUIs and CLIs	910
Accessing administration tools	911
Windows services used by IBM Spectrum Control	912
Frequently Asked Questions	912
Protocols and standards	913
Web Based Enterprise Management	913
Storage Management Initiative Specification	913
Service Location Protocol	914
Simple Network Management Protocol	914
Fibre Channel Methodologies of Interconnects	914
IBM Spectrum Control technical community	915
Accessibility features for IBM Spectrum Control	915
Troubleshooting and problem determination	916
Troubleshooting in IBM Spectrum Control	917
Introduction to troubleshooting	917
GUI troubleshooting	917
Renewing security certificates	917
Displaying performance metadata	919
Troubleshooting FTP transfers	919
The status of jobs is not updated after the Device server is stopped	921
Changing the session timeout value for the web-based GUI	921
Error message: This web browser is not supported	921
The GUI installation program displays the incorrect title on AIX 7.x with the UTF-8 locale	922
Error message: Out of memory at line: 14	922
Interface icons are not displayed in the correct positions	922
The status of a virtual machine is not consistent	923
Information exported to a CSV file is incorrectly formatted by spreadsheet program	923
General troubleshooting	923
Alerts aren't being generated	924
Performance information is not displayed for a resource	924
The Data server is shut down automatically	925
Fabric probe of DCFM CIM/OMs returns Java "Out of Memory" errors	925

Cluster resource group alerts are not triggered	925
Slow performance when moving data from the database repository	926
Error message when running repocopy on server that uses remote database	926
The Data Server service fails with a logon failure when restarted	926
Specifying a LUN ID for the assignvol command	926
Cannot monitor an EMC CIMOM on a Solaris server	927
SQLCODE-440 displayed if install IBM Spectrum Control on system with bad system clock	927
CIM agent runs slowly	927
Assigned LUN is not recognized by the host	927
On AIX systems, numbers in the Storage Resource agent registry are printed as formatted numbers	928
Fabric probe job failed	928
Error Messages: Get SQLCODE: -964, SQLSTATE: 57011, Unexpected error occurred while performing [a function]	929
Data server crashes on AIX	929
No performance data is retrieved	929
Performance monitoring job fails immediately after it starts	930
Error message "performance data files could not be correlated"	930
Performance correlation step takes a long time	930
Cannot define a performance monitor for a switch	930
File system storing database tables runs out of space	930
Error message: There is a problem with this website's security certificate	931
Cannot discover fabric on Red Hat Linux using Emulex HBA	931
Cannot get version of Storage Resource agent on SUSE Linux	932
Cannot log on to two instances of IBM Spectrum Control by using the same browser	932
Warning in Probe Log for Storage Resource agents on AIX (STA0115W)	932
A generic Storage Resource agent probe error might be a registration problem / GUID mismatch	933
Following IBM Spectrum Control server upgrade, the Storage Resource agents (SRA) might lack the Upgrade action	933
How to fix IBM Storwize V7000 Unified authentication errors	933
Slow performance on VMware virtual machines due to shared VMDK	934
How to disable cipher block chaining (CBC) ciphers in IBM Spectrum Control in response to the Lucky 13 vulnerability.	934
Db2 and database troubleshooting	935
Db2 crashes on RHEL 7.2	935
Db2 log messages ADM5530W can be ignored	936
Db2 log files fill the C drive on the Windows DB2 system	936
Cannot read Korean language logs from GUI	936
Installation fails with "No valid local database found" error	936
Error message: GUI0023M	936
Error message: SRV0024E	937
Error message: SRV0044E	937
Db2 instance crashes	937
Db2 connection errors	938
Error message: SQL0968C	938
Error message when installing Db2 on Windows	938
Error message: HWNOP0033 Database operation failed	939
Error message: DB2 SQL error: SQLCODE: -964, SQLSTATE: 57011	939
Error message: DB2 SQL error: SQLCODE: -973,STATE:57011, SQLERRMC:PCKCACHESZ	939
Error message: DEBUG: Error 2836	940
Error message: SRV0046E, Db2 stops running	940
Error message: The statement is too long or too complex (SQL0101N)	940
Db2 expired license error	941
Context root not found on IBM Spectrum Control 5.3.5 or later on Red Hat Enterprise Linux 7	941
Installation, uninstallation, and upgrading	942
Creating a keystore for IBM Spectrum Control server	942
Error message: INS3105E Failure to install remote agent	944
Installing a remote agent fails	944
The Solaris HBA identifies a SCSI device instead of fibre channel LUNs	944
Installing IBM Spectrum Control on Windows with installation files on a Samba fileshare	945
IBM Spectrum Control agent installation fails when files are mapped to a drive in a Windows Terminal Services environment	945
Db2 services do not start when you restart the system on a UNIX platform	945
Cannot install the Storage Resource agent on AIX through RSH	946
Storage Resource agent failed to install using install command	946
Upgrading to IBM Spectrum Control 5.4.1 on RHEL 7 fails during upgrade of Storage Resource agent	946
You cannot install IBM Spectrum Control on AIX and Linux systems	946
After reboot of the system, IBM Spectrum Control cannot communicate with Db2	947
Cannot deploy Storage Resource agent because of a timeout, get error message NAD0006E	947
Storage Resource agent information is inconsistent after stop and restart on Windows	948
Cannot deploy Storage Resource agent after changing SSL certificates	948
Message in Device server log states that port 162 is in use	949
Unable to log on to the GUI when the user name is defined in both the local OS repository and the Windows domain repository	949
BPCIN0057E and BPCIN0072E error messages: Installation on Linux operating system cannot proceed because of issues with DB2 db2inst1 user	950
Java virtual machine error during uninstall of IBM Spectrum Control	950
CAM-CRP-1613: Regenerating cryptographic keys in Cognos Analytics 11	950
BTACS0043E: Probe fails after upgrading IBM Spectrum Control	950

SAN Volume Controller troubleshooting	951
All ports of a host lose access to the volume	951
Error message: performance data files could not be correlated	951
SAN Volume Controller reports are incorrect	952
Probing a SAN Volume Controller cluster fails when the cluster ID changes	952
Setting the rate at which data is copied between volumes on storage systems that run IBM Spectrum Virtualize	952
Some capacity values are zero	953
DS8000 troubleshooting	953
Host connections for DS8000 storage systems are not being displayed	953
The user account is locked after the DS Storage Manager password for the HMC is changed	953
Get warning messages STA0044I, STA0035W, and STA0036I	954
Independent software vendors troubleshooting	954
HDS HiCommand shows a different set of volumes	955
Reports for an Hitachi Data Systems subsystem do not show current information	955
Tool to dump MSCS cluster configuration	955
Cannot collect performance data from EMC CLARiiON	955
Error message: HWNPM2132W	956
Assigning a volume for EMC CLARiiON fails	956
Incorrect password for SAN Volume Controller is accepted	956
You cannot unassign volumes from a Hewlett Packard EVA device	956
Getting timeouts and probe failures when using a Brocade switch	957
Resolving a problem connecting to ESX	957
Language troubleshooting	958
Changing the language in a web browser doesn't change the entire IBM Spectrum Control GUI	958
PDF files do not export correctly for some languages	958
Traditional Chinese characters not displayed correctly in GUI help panels	959
Cannot read Korean language logs from GUI	936
Resolving problems	961
IBM Spectrum Control tools	962
Repository copy tool	962
Exporting repository data	962
Service tool overview	963
Packaging log files from the command line and sending them to IBM Support	964
Creating a compressed file for a Storage Resource agent	965
How to customize the service tool	966
Tracing servers and agents	967
IBM Spectrum Control problem determination	970
Configuration files	970
server.config file	971
scheduler.config file	972
TPCD.config file	972
agent.config file	973
Log files	973
Installation log files for IBM Spectrum Control	973
Audit logs	974
Diagnosing IBM Spectrum Control problems	975
Discovery	976
Discovery and probe completion codes	977
Monitoring service	978
Performance monitoring	979
Configuration History	979
SMI-S fabric probe	979
VMware ESX	980
FlashCopy	981
tpctool	981
Fabric-specific problems	983
SMI-S providers	985
Linux SRA probes hanging problem	986
Getting support	986
Messages	987
Introduction to messages	987
Message types	988
AGT - Storage Resource Agent messages	989
AGT0001E Cannot set working directory to directory	992
AGT0002E Invalid invocation: swtchusr.	993
AGT0003I Agent registered.	993
AGT0004I Agent started.	993
AGT0005I Waiting for swtchusr.	993
AGT0006E Process process number is already connected to the server.	993
AGT0007E Another process is already connected to the server.	994
AGT0008E Job-type should be job type, but it is job type.	994
AGT0009E Error writing to swtchusr.	994

AGT0010E Class not found: class name.	994
AGT0011E class name does not inherit from class name.	994
AGT0012E class name does not have appropriate constructor.	995
AGT0013E Error constructing class name.	995
AGT0014E Cannot close the log-file log file name	995
AGT0015E The process identified by the process number failed to start generating the error status number.	995
AGT0016E Read from pipe failed.	996
AGT0017E Write to pipe failed.	996
AGT0018E Close of pipe failed.	996
AGT0019W Problem reading GUID identified by the GUID number.	996
AGT0020E Unparseable class-loader string class loader.	997
AGT0021E Error stopping the bundle.	997
AGT0031E User user name does not exist.	997
AGT0032E Cannot exec user command shell shell -- errno = arguments error message.	997
AGT0033E initgroups command(shell) fails -- errno = arguments error message.	998
AGT0034E Cannot create log-file in directory -- errno = arguments error message.	998
AGT0035W Error reading logical volume.	998
AGT0036W Error deserializing from file name.	998
AGT0037W Premature end of file -- file name.	998
AGT0038W Class class name not found restoring from file name.	998
AGT0039W Object restored from file name is not class name, Object is class name.	999
AGT0040E Agent Shutting down.	999
AGT0041E Agent not registered.	999
AGT0042E Error writing file name.	999
AGT0043E Error serializing to file name.	999
AGT0044E Fatal error -- cannot connect to self port number.	1000
AGT0045E Error parsing configuration file on line line number.	1000
AGT0046W In agent.config, key name is not an integer -- default of default key number used.	1000
AGT0047E In agent.config, key name is not a valid port.	1000
AGT0048E In agent.config, no value supplied for key name.	1000
AGT0049E Too many unreported jobs.	1001
AGT0050E Error parsing configuration file.	1001
AGT0051W The configuration file had at least one bad value.	1001
AGT0052W Cannot create temporary file in directory.	1001
AGT0053W Tried to send signal process number to nonexistent PID process number.	1001
AGT0054W No running job (job name, job number).	1001
AGT0055E Jobs are not allowed to run as root.	1002
AGT0056I NOTICE: Server moved to computer name:port number.	1002
AGT0057E Unable to create message-queue from key value.	1002
AGT0058E Error waiting for job to start.	1002
AGT0059I Received request to shut down request type.	1002
AGT0060I Exiting normally.	1002
AGT0061I Exiting abnormally.	1003
AGT0062E Cannot fork job.	1003
AGT0063E Error terminating process.	1003
AGT0064E Putative SID has too few hyphens: security identifier.	1003
AGT0065E Error waiting for shutdown request with Queue ID.	1003
AGT0066E Error reaping.	1004
AGT0067E Error waiting for process process name.	1004
AGT0068E Cannot fork -- command: command name.	1004
AGT0069E open(/proc/uptime) fails.	1004
AGT0070E read(/proc/uptime) fails.	1005
AGT0071E /proc/uptime does not look right.	1005
AGT0072E Cannot determine system-boot time.	1005
AGT0073E Problem creating up-time poller.	1005
AGT0074E Unsupported Operating System.	1006
AGT0075E Cannot query key name key value.	1006
AGT0076E Performance title data title not found.	1006
AGT0077E <<System>> performance object not found.	1006
AGT0078E The <<system>> performance counter not found.	1006
AGT0079E Cannot create semaphore.	1007
AGT0080W Cannot retrieve process times.	1007
AGT0081W Cannot retrieve process exit code.	1007
AGT0082E Cannot log in user user name.	1007
AGT0083E Cannot obtain SID of local computer (computer name).	1008
AGT0084E Cannot get name of local computer.	1008
AGT0085E SID of local computer has unexpected type (security identifier type).	1008
AGT0086W The GetTokenInformation method failed to execute.	1008
AGT0087W The LookupAccountSid method failed to execute.	1008
AGT0088W Cannot find any domain controller for domain domain name.	1009
AGT0089W Cannot fetch the information for user domain\\user.	1009
AGT0090W Cannot load profile for user username.	1009

AGT0091W Cannot unload user profile.	1009
AGT0092E Unable to retrieve status of job with the job number.	1010
AGT0093E Cannot retrieve environment block.	1010
AGT0094E Cannot impersonate to create resource.	1010
AGT0095W Environment for user username not fully configured.	1010
AGT0096E Malformed command-line.	1011
AGT0097E Command not found: fdisk command.	1011
AGT0098E Command has unsupported extension.	1011
AGT0099E Cannot duplicate handle.	1011
AGT0101E Cannot determine whether this is a domain controller.	1011
AGT0102E Cannot determine name of own domain.	1012
AGT0103W Cannot make Windows Job Object.	1012
AGT0104E Wrong server.	1012
AGT0105W Ignoring server relocation in config file.	1012
AGT0106I Server has changed.	1013
AGT0110E Cannot open key key name key value.	1013
AGT0111I Rereading config file.	1013
AGT0112E This product is not fully installed. To try again, stop and restart this agent.	1013
AGT0113E Cannot create temporary file in directory error message.	1013
AGT0114E Cannot write to file file name error message.	1013
AGT0115E Fork failed.	1014
AGT0116E Cannot exec command name.	1014
AGT0117W Cannot open file for auto-delete.	1014
AGT0118E Upgrader is not okay.	1014
AGT0119I version, modification and release.	1015
AGT0120E Error transmitting shutdown request to agent.	1015
AGT0121E Unable to create socket.	1015
AGT0122I The agent is already down.	1015
AGT0123W Unable to determine if agent is active. Sending shutdown request.	1015
AGT0124E Unable to create restartable job directory: directory name.	1016
AGT0125E Unable to read directory directory name.	1016
AGT0126E Unable to restart job from file file name.	1016
AGT0127E Unable to restart job job name.name run number run number.	1016
AGT0128W Cannot retrieve global structure.	1017
AGT0129W A system call failed in the agent program.	1017
AGT0130E The upgrader program parameters are missing.	1017
AGT0131I Exit Status = exit message.	1017
AGT0132E Cannot find/create script <script name>.	1017
AGT0133I Running Command: script name script contents.	1017
AGT0134E Putative SID does not start correctly: security identifier.	1017
AGT0135E Component component name is not a number: component number.	1018
AGT0136E No INSTANCEn.DAT file found!	1018
AGT0137E Cannot determine working directory!	1018
AGT0138W product name is not licensed on this computer.	1018
AGT0139E An initial probe cannot be performed.	1019
AGT0140W Discovery will not be performed.	1019
AGT0141W A filesystem scan will not be performed.	1019
AGT0142E Cannot make handle inheritable.	1019
AGT0143I file_count files scanned	1019
AGT0144I file_count total files scanned	1019
AGT0145I Retrieving job definition from server	1019
AGT0146I Scan started	1020
AGT0147I Retrieving report partition partition_number	1020
AGT0148I Report data retrieved	1020
AGT0149I Retrieving history data	1020
AGT0150I Deleting temporary files	1020
AGT0151E Unable to retrieve report definition	1020
AGT0152I Job definition retrieved	1020
AGT0153E MSCSEventListener thread interrupted. Shutting down native MSCS event listener.	1021
AGT0154E No IP address could be found for the local host.	1021
AGT0155I Windows Scan Option : scan parameters.	1021
AGT0156W Error killing process process id.	1021
AGT0157E Interrupted (agent shutting down).	1021
AGT0159E A problem was encountered stopping IBM Spectrum Control subagent.	1022
AGT0160E Error enumerating keys under registry path.	1022
AGT0161W Error querying value root path\\subkey name\\library file.	1022
AGT0162E Library : HBA_RegisterLibrary returned value.	1022
AGT0163E Library : HBA_LoadLibrary returned value.	1023
AGT0164W Library : HBA_GetAdapterName(adapter number) returned adapter name.	1023
AGT0165W Library : HBA_OpenAdapter(adapter name) failed.	1023
AGT0166W The HBA API HBA_GetAdapterAttributes for adapter adapter name returned error error code.	1023
AGT0167W Library : HBA_GetAdapterPortAttributes(adapter number, port) returned status.	1024

AGT0168W HBA data cannot be collected because the adapter name> adapter on the server does not support the HBA_GetAdapterAttributes function.	1024
AGT0200E Error waiting for process.	1024
AGT0250E Error starting bundle: exception message.	1024
AGT0251E Failed to install language pack.	1024
AGT0252E Failed to install language pack: File not found archive	1025
AGT0253E Invalid file format: file name	1025
AGT0254E Failed to load language pack.	1025
AGT0256I Waiting for Common Agent services.	1025
AGT0259E Agent cannot upgrade, the required space of 50 meg was not met.	1025
AGT0260E Failed to create a session with Reliable Scalable Cluster Technology Error ID-Error type: Error Description	1026
AGT0261E Failed to end a session with Reliable Scalable Cluster Technology Error ID-Error type: Error Description	1026
AGT0262W Concurrent cluster resource groups are not supported. The cluster will not be probed.	1026
AGT0263E Failed to get resource data from the HACMP cluster using the RSCT RMC interface.	1026
AGT0264E Failed to determine if the local node is clustered using the following command: command line	1027
AGT0265E Failed to get the cluster name using the following command: command line	1027
AGT0266E Failed to get the cluster ID using the following command: command line	1027
AGT0267E Failed to get the names of all the resource groups in the cluster.	1027
AGT0268E Failed to get the list of all the resource groups.	1028
AGT0269E Failed to get the service IP label information for resource group cluster resource group.	1028
AGT0270E Failed to get the volume group resources associated with cluster resource group cluster resource group.	1028
AGT0271I query command command selected.	1028
AGT0271E Failed to get the export resources associated with cluster resource group cluster resource group.	1029
AGT0272E Failed to get state of all the resource groups in the cluster.	1029
AGT0273E Failed to query the ODM query command.	1029
AGT0274E Storage Resource Agent initialization failed, return code: return code.	1029
AGT0275E Failed to get the HACMP node name using the following command:\ \n command line	1030
AGT0276E Failed to get the physical volume information for the volume group volume group.	1030
AGT0277E Failed to get the logical volume information for the volume group volume group.	1030
AGT0278E Failed to get the SDD device information using the following command: Command	1030
AGT0279W The detected level of HACMP is not compatible. Please consult the user's guide for compatible versions.	1031
AGT0279I Registry entry added for installation location location.	1031
AGT0280W Failed to determine if the cluster is stable using the following command.	1031
AGT0281I Config file entry added for server server and port number number.	1031
AGT0281W The cluster is not ready to be probed. Waiting settle time seconds to retry (retry count/retries).	1031
AGT0282E Unable to probe the cluster.	1032
AGT0283I The cluster is ready and the probe will proceed.	1032
AGT0284E Agent host name does not accept scripts from server (scripts are "Disabled"). Script script name cannot run.	1032
AGT0285E The script script name has an incorrect Windows extension.The accepted Windows extensions are: extension list.	1032
AGT0301I Extracting file	1033
AGT0304E Failed to send agent registration for agent in server to server server.	1033
AGT0305I Successfully sent Agent, server, registration to server number	1033
AGT0306E Unable to send status to server.	1033
AGT0307I Successfully sent Probe complete status to server, server.	1033
AGT0314E Unable to stop all jobs for server server name	1034
AGT0383I Install completed successfully.	1034
AGT0389E Installation aborted. Communication method specified for this installation is communication method and does not match communication method, the communication method of the installed SRA.	1034
AGT0404I Creating Common Agent Package file file name	1034
AGT0405I Creating Common Agent Package directory directory name	1034
AGT0406I Common Agent Package file file name	1034
AGT0407I Common Agent Package directory directory name	1034
AGT0408E Failed to create Common Agent Package directory directory name	1035
AGT0409E Failed to create Common Agent Package file file name	1035
AGT0410E Could not delete Common Agent Package file file name	1035
AGT0411E SRAutil: invalid command command name	1035
AGT0412E SRAutil: invalid option options name	1036
AGT0413E Unable to start script script name	1036
AGT0414I Script script name	1036
AGT0415I Unable to load HBA library, rc: return code	1036
AGT0416I Number of HBA adapters on the system: number of adapters	1036
AGT0417I Adapter adapter name	1036
AGT0418I Found switched fabric: fabric ID	1036
AGT0419I Interconnected element element ID	1037
AGT0420I There are no adapters connected to switch fabric	1037
AGT0421E Unable to create fabric fabric component data file: data file	1037
AGT0422E Error writing to fabric fabric component data file: data file	1037
AGT0423I fabric command command on switch fabric failed, reason: reason code, explanation: explanation code	1037
AGT0424I Storage Resource Agent was unable to retrieve the fabric name for switch fabric fabric ID, reason: reason code, explanation: explanation code	1037
AGT0426I Fabric Probe Data: data field	1038
AGT0427I Fabric Discovery Data: data field	1038
AGT0428I Adapter adapter index not connected to switch fabric.	1038
AGT0429I Process process name invoked with command line arguments command arguments	1038

AGT0430I Process process name exiting with return code return code _____	1038
AGT0431W HBA API call function name failed with return code t _____	1038
AGT0432E Insufficient response buffer size passed for command command name _____	1038
AGT0433E response for command command name contains invalid data _____	1039
AGT0434E Errors parsing command name _____	1039
AGT0435I Command : command name _____	1039
AGT0436I Command : command arguments _____	1039
AGT0437I Error error code on adapter adapter index, unable to determine connection to switch fabric. _____	1039
AGT0438I Response: data field _____	1040
AGT0439E Errors writing command name _____	1040
AGT0440E Zone control command execution failed _____	1040
AGT0441E Unable to verify command completion, rc: return code _____	1040
AGT0442I ----- BEGIN OUTPUT ----- _____	1040
AGT0443I ----- END OUTPUT ----- _____	1041
AGT0444I Switched fabric: fabric ID already found _____	1041
AGT0445I Operation has been cancelled _____	1041
AGT0446I Fabric discovery successfully found switch fabric. _____	1041
AGT0447I Fabric discovery did not find switch fabric. _____	1041
AGT0448I Fabric discovery is already running. _____	1041
AGT0449I Fabric discovery failed, unable to discover switch fabric. _____	1041
AGT0452E No Multipath Device mapped to ID id _____	1041
AGT0453E Error when setting policy on device device: error _____	1042
AGT0454I For Multipath DM driver, only Round Robin policy is available. _____	1042
AGT0455W Setting policy for Multipath EMC Powerpath driver is not supported. _____	1042
AGT0456E No supported multipath driver was found on this system. _____	1042
AGT0457E Multipath policy configuration is not supported for this multipath driver. _____	1042
AGT0458I Setting multipath policy policy for following multipath devices: devices _____	1043
AGT0459I Executing Walk the Bus action for refreshing system configuration. _____	1043
AGT0460E Invalid source or target directory (directory). Source directory must be the installation image location. Target directory cannot be the same as source directory. _____	1043
AGT0461E Installation path contains an invalid character for the target platform: character _____	1043
AGT0462E At least one directory component in install location contains a reserved name for the target platform: name _____	1043
AGT0463I DM Multipath driver is installed but not loaded. _____	1044
AGT0464I multipath.conf could not be found on the system. _____	1044
AGT0469I Not enough disk space on disk. At least sizeMB is required! _____	1044
AGT0484E Cannot obtain the hostname of the Storage Resource Agent. _____	1044
AGT0485I The Storage Resource Agent on the IBM Spectrum Control server cannot be deleted. _____	1044
AGT0486W The available disk space is low on disk. It is recommended that at least size MB of disk space is available on the disk partition for successful operation of the Storage Resource agent. _____	1045
AGT0487E Not enough disk space available on disk. At least 20 MB of available disk space is required to complete the probe. Increase the available disk space on the disk partition and then start the probe again. _____	1045
AGT0504I Validating user user name _____	1045
AGT0505I User user name validation succeeded. _____	1045
AGT0506E User user name validation failed. _____	1045
AGT0507W User user name does not exist, user will be created. _____	1045
AGT0509I User user name created successfully. _____	1046
AGT0510E Failed to create user user name _____	1046
AGT0511E Usage error: -duser is missing. _____	1046
AGT0512E Usage error: -dpassword option is missing. _____	1046
AGT0513I Successfully sent Scan complete status to server, agent install directory _____	1046
AGT0514E Failed to send Scan complete status to server, agent install directory _____	1047
AGT0515I Stopping all SRA jobs ... _____	1047
ALR - Spectrum Control Alert messages _____	1047
ALR0001I The amount of RAM on host server name has changed from current value to new value. _____	1057
ALR0002I The amount of virtual memory on host server name has changed from current value to new value. _____	1057
ALR0003I A new disk drive has been detected on host server name. Disk manufacturer/serial Number: manufacturer/serial Number. _____	1057
ALR0004E A previously visible disk drive can no longer be found on host server name. Disk manufacturer/serial number: manufacturer/serial number. _____	1057
ALR0005I A new filesystem has been detected on host server name. Filesystem mount point: mount point. _____	1057
ALR0006E A previously visible filesystem can no longer be found on host server name. Filesystem mount point: mount point. _____	1057
ALR0007E A disk drive visible on host server name has predicted that a disk failure is imminent. Disk manufacturer/serial number: manufacturer/serial number. _____	1058
ALR0008I The physical space definition of filesystem filesystem name on host server name has been reconfigured. _____	1058
ALR0009W The free space on filesystem filesystem name on host server name has fallen below the threshold value of threshold. The free space is freespace or percent of the filesystem capacity. _____	1058
ALR0010W The number of free inodes on filesystem filesystem name on host server name has fallen below the threshold value of threshold. The number of free inodes is free inodes or percent of the filesystem's total inodes. _____	1058
ALR0011W A new grown defect has been detected on a disk visible to host server name. Disk manufacturer/serial Number: disk manufacturer/serial number, Current grown defects: current, Previous grown defects: previous. _____	1058
ALR0012W The number of grown disk defects has exceeded the threshold value of threshold. Host: server name, Disk manufacturer/serial number: disk manufacturer/serial number, Current grown defects: current, Previous grown defects: previous. _____	1058
ALR0013W A new monitored directory has been detected on host server name. Directory name: directory name, Directory Group: directory group. _____	1059
ALR0014E A monitored directory has been removed from host server name. Directory name: directory name, Directory group: directory group. _____	1059
ALR0015W Directory directory on host server name has exceeded its space usage quota of quota. The directory is currently consuming usage or percent of the filesystem capacity. _____	1059

ALR0016W Filesystem filesystem name on host server name has violated a filesystem constraint. number of files file(s) consuming space or percent of the filesystem capacity are in violation of the conditions defined in this constraint. The constraint threshold is threshold. User user name} has number of files} files consuming space} of storage. violating owners}.	1059
ALR0017E Host server name appears to be down. number of attempts attempt(s) to ping this host have failed.	1059
ALR0018W quota name user has exceeded a network storage usage quota of usage. This user is currently consuming amount of storage.	1059
ALR0019W quota name user on host server name has exceeded a server storage usage quota of usage. This user is currently consuming amount of storage.	1060
ALR0020W quota name user on host server name has exceeded a filesystem usage quota of <usage>valueon filesystem filesystem name. This user is currently consuming amount of storage.	1060
ALR0021W Run number number of job creator job name has failed on run number of total jobs total jobs.	1060
ALR0022I Server server name has been discovered.	1060
ALR0023W Run number run number of job creator job name has failed.	1060
ALR0024W User user name on host server name has exceeded a table space usage quota of usage on table space rdbms type, table space instance. This user is currently consuming amount of storage.	1060
ALR0025W User user name on host server name has exceeded an RDBMS instance usage quota of quota on rdbms type instance instance name. This user is currently consuming amount of storage.	1061
ALR0026W User user name has exceeded a network database storage usage quota of quota. This user is currently consuming amount of storage.	1061
ALR0027W The log directory file name archived log directory value on host server name has exceeded the threshold value of threshold. This directory currently contains number of logs archived logs consuming amount of storage.	1061
ALR0028I A new table space name has been discovered on rdbms type instance on host server name. rdbms instance: database.	1061
ALR0029E value value has been dropped. RDBMS: value value, Host: host name.	1061
ALR0030E value value has been taken offline. RDBMS: value value, Host: host name.	1061
ALR0031W The free space on value: value, RDBMS: value value, host: host name, has fallen below the threshold value of value. The free space is value or value of the value capacity.	1061
ALR0032W The free space on table space: table space name, RDBMS: rdbms type instance, database: database, host: server name, is fragmented across number of extents extents. This exceeds the threshold value of threshold extents. The largest contiguous free extent is largest extent.	1062
ALR0033W The largest free extent available on table space: table space name, RDBMS: rdbms type rdbms name, database: database name, host: server name, has fallen below the threshold value of threshold. The largest free extent is largest extent.	1062
ALR0034W Segment segment name of table/cluster table name on host: server name, RDBMS: rdbms type rdbms name, database: database name, is fragmented across number of extents extents. This exceeds the threshold value of threshold extents. This segment is a value type segment.	1062
ALR0035W Segment value of table/cluster value on host: host name, RDBMS: value value, database: value, is nearing the maximum number of extents available to it. This segment currently occupies value extents. The value additional extent(s) available to this segment falls below the defined threshold of value extent(s). The segment is a type type segment.	1062
ALR0036W Table/cluster table name on host: server name, RDBMS: rdbms type rdbms name, database: database name, has exceeded a space usage quota of usage quota. This table is currently consuming amount of storage.	1062
ALR0037W Table/cluster value on host: host name, RDBMS: value value, database: database name, has exceeded a chained row quota of value. Statistics indicate that value rows or value of the total rows are chained.	1063
ALR0038W Segment name of table/cluster table name on host: server name, RDBMS: rdbms type rdbms name, database: database name, has amount of unused, wasted space. This represents amount of the total space allocated to the segment, and exceeds the threshold value of threshold. This segment is a type type segment.	1063
ALR0039E Table/cluster table name on host: server name, RDBMS: rdbms type rdbms name, database: database name, has been dropped.	1063
ALR0040W Filer filer name has been discovered. Spectrum Control will not monitor this filer until it has been licensed.	1063
ALR0041W The amount of log freespace available on instance: database, RDBMS: rdbms type , <rdbms>host: server name, has fallen below the threshold value of threshold. The amount of log freespace available is amount or percent of the total capacity.	1063
ALR0042I A new device has been discovered on rdbms instance instance on host server name. Device: device name, Capacity: capacity, File Name: filename.	1063
ALR0043E Device device name has been dropped from rdbms instance instance on host server name. Capacity: capacity, File Name: filename.	1064
ALR0044W The amount of device freespace available on value instance value on host value has fallen below the threshold value of value. The amount of device freespace available is value or value of the current capacity of value.	1064
ALR0045W The amount of device freespace available on device instance instance on host server name has gone above the threshold value of threshold. The amount of device freespace available is freespace or percent of the current capacity of total capacity.	1064
ALR0046W Database database name has not been backed up in the last number of days days. Last backup for the database was on date. RDBMS: rdbms type rdbms name, Host: server name.	1064
ALR0047W filer name filer type is no longer accessible from host entity.	1064
ALR0048W Storage Subsystem subsystem name is no longer accessible from host server name.	1064
ALR0049W disk array name disk array type has been discovered from host server name. Spectrum Control will not monitor this disk array until it has been selected as for monitoring from within the Storage Subsystem Administration GUI.	1064
ALR0050W The amount of cache on storage subsystem subsystem name has changed from old value to new value.	1065
ALR0052W Filesystem filesystem name on host server name will be automatically extended because its free space has fallen below the threshold of threshold. Current free space: current free space; Current capacity: current capacity; Target capacity: target capacity}.	1065
ALR0053W Filesystem filesystem name on host server name needs extension but will not be because its current capacity of current capacity exceeds the specified limit of limit. Filesystem free space: freespace (current capacity of current capacity).	1065
ALR0055I Cluster resource group cluster resource group name was added to cluster cluster name on node node name .	1065
ALR0056I Cluster resource group cluster resource group name was removed from cluster cluster name on node node name .	1065
ALR0057I Cluster resource group cluster resource group name was moved in cluster cluster name from node node name to node node name.	1065
ALR0076W Performance monitor failure for device value.	1066
ALR0500E The Disk Utilization Percentage of array array name in storage system storage system name was measured to be measured value%. This violated the critical-stress boundary value of boundary value%.	1066
ALR0501W The Disk Utilization Percentage of array array name in storage system storage system name was measured to be measured value%. This violated the warning-stress boundary value of boundary value%.	1066
ALR0502W The Disk Utilization Percentage of array array name in storage system storage system name was measured to be measured value%. This violated the warning-idle boundary value of boundary value%.	1067
ALR0503E The Disk Utilization Percentage of array array name in storage system storage system name was measured to be measured value%. This violated the critical-idle boundary value of boundary value%.	1067
ALR0504E The Total Back-end I/O Rate of array, MDisk, or MDisk Group Array, MDisk, or MDisk Group name in storage system storage system name was measured to be measured value ops/s, which violated the defined critical-stress boundary value of boundary value ops/s.	1067

[illegible]

[illegible]

ALR0579E The Port to Local Node Send Response Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-idle boundary value of boundary value ms/op. 1094

ALR0580E The Port to Local Node Receive Response Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-stress boundary value of boundary value ms/op. Violation of this threshold boundary could mean that it is taking too long to send data between nodes on the fabric, which suggests either a problem with the nodes or congestion around the associated FC ports on the fabric. 1094

ALR0581W The Port to Local Node Receive Response Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined warning-stress boundary value of boundary value ms/op. Violation of this threshold boundary could mean that it is taking too long to send data between nodes on the fabric, which suggests either a problem with the nodes or congestion around the associated FC ports on the fabric. 1095

ALR0582W The Port to Local Node Receive Response Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined warning-idle boundary value of boundary value ms/op. 1095

ALR0583E The Port to Local Node Receive Response Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-idle boundary value of boundary value ms/op. 1096

ALR0584E The Port to Local Node Send Queue Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-stress boundary value of boundary value ms/op. Violation of this threshold boundary could mean that the node has to wait too long to send data to other nodes on the fabric, which suggests congestion around the associated FC ports on the fabric. 1096

ALR0585W The Port to Local Node Send Queue Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined warning-stress boundary value of boundary value ms/op. Violation of this threshold boundary could mean that the node has to wait too long to send data to other nodes on the fabric, which suggests congestion around the associated FC ports on the fabric. 1096

ALR0586W The Port to Local Node Send Queue Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined warning-idle boundary value of boundary value ms/op. 1097

ALR0587E The Port to Local Node Send Queue Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-idle boundary value of boundary value ms/op. 1097

ALR0588E The Port to Local Node Receive Queue Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-stress boundary value of boundary value ms/op. Violation of this threshold boundary could mean that the node has to wait too long to receive data to other nodes on the fabric, which suggests congestion around the associated FC ports on the fabric. 1097

ALR0589W The Port to Local Node Receive Queue Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined warning-stress boundary value of boundary value ms/op. Violation of this threshold boundary could mean that the node has to wait too long to receive data to other nodes on the fabric, which suggests congestion around the associated FC ports on the fabric. 1098

ALR0590W The Port to Local Node Receive Queue Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined warning-idle boundary value of boundary value ms/op. 1098

ALR0591E The Port to Local Node Receive Queue Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-idle boundary value of boundary value ms/op. 1099

ALR0592E The Non-preferred Node Usage Percentage of I/O Group I/O Group name in device device name was measured to be measured value%, which violated the critical-stress boundary value of boundary value%. 1099

ALR0593W The Non-preferred Node Usage Percentage of I/O Group I/O Group name in device device name was measured to be measured value%, which violated the warning-stress boundary value of boundary value%. 1099

ALR0594W The Non-preferred Node Usage Percentage of I/O Group I/O Group name in device device name was measured to be measured value%, which violated the warning-idle boundary value of boundary value%. 1100

ALR0595E The Non-preferred Node Usage Percentage of I/O Group I/O Group name in device device name was measured to be measured value%, which violated the critical-idle boundary value of boundary value%. 1100

ALR0596E The Peak Back-end Write Response Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-stress boundary value of boundary value ms/op. If writes to disk are too slow, and writes are being received faster than they can be destaged to disk, then the nodes write-cache will eventually fill up. In extreme cases, the node will stop caching write data, causing a significant performance degradation for the affected volumes. 1100

ALR0597W The Peak Back-end Write Response Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined warning-stress boundary value of boundary value ms/op. If writes to disk are too slow, and writes are being received faster than they can be destaged to disk, then the nodes write-cache will eventually fill up. In extreme cases, the node will stop caching write data, causing a significant performance degradation for the affected volumes. 1101

ALR0598W The Peak Back-end Write Response Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined warning-idle boundary value of boundary value ms/op. 1101

ALR0599E The Peak Back-end Write Response Time of node node name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-idle boundary value of boundary value ms/op. 1102

ALR0600E The Port Send Utilization Percentage of port port name in storage system storage system name was measured to be measured value%, which violated the critical-stress boundary value of boundary value%. 1102

ALR0601W The Port Send Utilization Percentage of port port name in storage system storage system name was measured to be measured value%, which violated the warning-stress boundary value of boundary value%. 1103

ALR0602W The Port Send Utilization Percentage of port port name in storage system storage system name was measured to be measured value%, which violated the warning-idle boundary value of boundary value%. 1103

ALR0603E The Port Send Utilization Percentage of port port name in storage system storage system name was measured to be measured value%, which violated the critical-idle boundary value of boundary value%. 1103

ALR0604E The Port Receive Utilization Percentage of port port name in storage system storage system name was measured to be measured value%, which violated the critical-stress boundary value of boundary value%. 1104

ALR0605W The Port Receive Utilization Percentage of port port name in storage system storage system name was measured to be measured value%, which violated the warning-stress boundary value of boundary value%. 1104

ALR0606W The Port Receive Utilization Percentage of port port name in storage system storage system name was measured to be measured value%, which violated the warning-idle boundary value of boundary value%. 1105

ALR0607E The Port Receive Utilization Percentage of port port name in storage system storage system name was measured to be measured value%, which violated the critical-idle boundary value of boundary value%. 1105

ALR0608E The Port Send Bandwidth Percentage of port port name in device device name was measured to be measured value%, which violated the critical-stress boundary value of boundary value%. 1105

ALR0609W The Port Send Bandwidth Percentage of port port name in device device name was measured to be measured value%, which violated the warning-stress boundary value of boundary value%. 1106

ALR0610W The Port Send Bandwidth Percentage of port port name in device device name was measured to be measured value%, which violated the warning-idle boundary value of boundary value%. 1106

ALR0611E The Port Send Bandwidth Percentage of port port name in device device name was measured to be measured value%, which violated the critical-idle boundary value of boundary value%. 1106

[illegible]

[illegible]

ALR0686W The Port Congestion Index of port port_name in device device_name was found to be measured_value Counts, which violated the defined warning-idle boundary value of boundary_value Counts.	1132
ALR0687E The Port Congestion Index of port port_name in device device_name was found to be measured_value Counts, which violated the defined critical-idle boundary value of boundary_value Counts.	1132
ALR0688E The Link Quality Percentage of port port_name in device device_name was found to be measured_value, which violated the defined critical-stress boundary value of boundary_value.	1133
ALR0689W The Link Quality Percentage of port port_name in device device_name was found to be measured_value, which violated the defined warning-stress boundary value of boundary_value.	1133
ALR0690W The Link Quality Percentage of port port_name in device device_name was found to be measured_value, which violated the defined warning-idle boundary value of boundary_value.	1133
ALR0691E The Link Quality Percentage of port port_name in device device_name was found to be measured_value, which violated the defined critical-idle boundary value of boundary_value.	1134
ALR0692E The Invalid Link Transmission Rate of port port_name in device device_name was found to be measured_value, which violated the defined critical-stress boundary value of boundary_value.	1134
ALR0693W The Invalid Link Transmission Rate of port port_name in device device_name was found to be measured_value, which violated the defined warning-stress boundary value of boundary_value.	1134
ALR0694W The Invalid Link Transmission Rate of port port_name in device device_name was found to be measured_value, which violated the defined warning-idle boundary value of boundary_value.	1135
ALR0695E The Invalid Link Transmission Rate of port port_name in device device_name was found to be measured_value, which violated the defined critical-idle boundary value of boundary_value.	1135
ALR0696E The Extreme I/O Concurrency Percentage of port port_name in storage system device_name was found to be measured value, which violated the defined critical-stress boundary value of boundary value.	1135
ALR0697W The Extreme I/O Concurrency Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-stress boundary value of boundary_value.	1136
ALR0698W The Extreme I/O Concurrency Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-idle boundary value of boundary_value.	1136
ALR0699E The Extreme I/O Concurrency Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined critical-idle boundary value of boundary_value.	1137
ALR0700E The I/O Busy Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined critical-stress boundary value of boundary_value.	1137
ALR0701W The I/O Busy Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-stress boundary value of boundary_value.	1137
ALR0702W The I/O Busy Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-idle boundary value of boundary_value.	1138
ALR0703E The I/O Busy Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined critical-idle boundary value of boundary_value.	1138
ALR0704E The I/O Overrun Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined critical-stress boundary value of boundary_value.	1139
ALR0705W The I/O Overrun Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-stress boundary value of boundary_value.	1139
ALR0706W The I/O Overrun Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-idle boundary value of boundary_value.	1139
ALR0707E The I/O Overrun Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined critical-idle boundary value of boundary_value.	1140
ALR0708E The Zero Send Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined critical-stress boundary value of boundary_value.	1140
ALR0709W The Zero Send Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-stress boundary value of boundary_value.	1140
ALR0710W The Zero Send Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-idle boundary value of boundary_value.	1141
ALR0711E The Zero Send Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined critical-idle boundary value of boundary_value.	1141
ALR0712E The Zero Receive Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined critical-stress boundary value of boundary_value.	1142
ALR0713W The Zero Receive Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-stress boundary value of boundary_value.	1142
ALR0714W The Zero Receive Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-idle boundary value of boundary_value.	1142
ALR0715E The Zero Receive Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined critical-idle boundary value of boundary_value.	1143
ALR0716E The number of bytes received on each of the client network interfaces of cluster cluster name in storage subsystem storage system name was measured to be measured valuebytes, which violated the defined critical-idle boundary value of boundary valuebytes.	1143
ALR1114I New Storage Subsystem discovered.	1143
ALR4000I Endpoint device endpoint device name has been discovered.	1143
ALR4001W Endpoint device endpoint device name is missing.	1144
ALR4002I Endpoint device endpoint device name has been rediscovered.	1144
ALR4015I subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN has been discovered.	1144
ALR4016W subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN is missing.	1144
ALR4017I subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN has been rediscovered.	1144
ALR4018W subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN has gone offline.	1145
ALR4019I subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN has gone online.	1145
ALR4020I Switch switch name or WWN has been discovered.	1145
ALR4021E Switch switch name or WWN is missing.	1145
ALR4022I Switch switch name or WWN has been rediscovered.	1145
ALR4023W The version for switch switch name or WWN has changed from previous firmware version to current firmware version .	1146

ALR4024W Status of switch switch name or WWN has degraded from previous status to current status .	1146
ALR4025I Status of switch switch name or WWN has improved from previous status to current status .	1146
ALR4026I Blade blade name or WWN on switch switch name or WWN has been discovered.	1146
ALR4027W Blade blade name or WWN on switch switch name or WWN is missing.	1146
ALR4028I Blade blade name or WWN on switch switch name or WWN has been rediscovered.	1147
ALR4029E Blade blade name or WWN on switch switch name or WWN has gone offline.	1147
ALR4030I Blade blade name or WWN on switch switch name or WWN has gone online.	1147
ALR4034W The driver version for HBA adapter name on server host name has changed from previous version to new version.	1147
ALR4035W The firmware version for HBA adapter name on server host name has changed from previous version to new version .	1147
ALR4046I Fabric fabric name or WWN is discovered.	1148
ALR4047E Fabric fabric name or WWN is missing.	1148
ALR4048I Fabric fabric name or WWN is rediscovered.	1148
ALR4051I Inactive zone zone name in fabric fabric name or WWN has been discovered.	1148
ALR4052W Inactive zone zone name in fabric fabric name or WWN is missing.	1148
ALR4053I Inactive zone zone name in fabric fabric name or WWN has been rediscovered.	1148
ALR4054I Inactive zoneset zoneset name in fabric fabric name or WWN has been discovered.	1149
ALR4055W Inactive zoneset zoneset name in fabric fabric name or WWN is missing.	1149
ALR4056I Inactive zoneset zoneset name in fabric fabric name or WWN has been rediscovered.	1149
ALR4063I The connection from switch or node switch name or WWN port port name or WWPN to switch or node switch name or WWN port port name or WWPN has been discovered.	1149
ALR4064W The connection from switch or node switch name or WWN port port name or WWPN to switch or node switch name or WWN port port name or WWPN is missing.	1149
ALR4065I The connection from switch or node switch name or WWN port port name or WWPN to switch or node switch name or WWN port port name or WWPN has been rediscovered.	1150
ALR4066I Switch switch name or WWN in fabric fabric name or WWN has been discovered.	1150
ALR4067W Switch switch name or WWN in fabric fabric name or WWN is missing.	1150
ALR4068I Switch switch name or WWN in fabric fabric name or WWN has been rediscovered.	1150
ALR4069I Port port name or WWPN in switch switch name or WWN has been discovered.	1150
ALR4070W Port port name or WWPN in switch switch name or WWN is missing.	1151
ALR4071I Port port name or WWPN in switch switch name or WWN has been rediscovered.	1151
ALR4078I Alias zone alias has been added to inactive zone zone name in fabric fabric name or WWN .	1151
ALR4079W Alias zone alias has been removed from inactive zone zone name in fabric fabric name or WWN .	1151
ALR4080I Alias zone alias has been readded to inactive zone zone name in fabric fabric name or WWN .	1151
ALR4081I Zone member zone member name has been added to inactive zone zone name in fabric fabric name or WWN .	1152
ALR4082I Zone member zone member name has been removed from inactive zone zone name in fabric fabric name or WWN .	1152
ALR4083I Zone member zone member name has been readded to inactive zone zone name in fabric fabric name or WWN .	1152
ALR4084I Zone zone name has been added to inactive zone set zone set name in fabric fabric name or WWN .	1152
ALR4085I Zone zone name has been removed from inactive zone set zone set name in fabric fabric name or WWN .	1152
ALR4086I Zone zone name has been readded to inactive zone set zone set name in fabric fabric name or WWN .	1153
ALR4089W ZoneSet zoneset name in fabric fabric name or WWN has been deactivated. ZoneSet zoneset name has been activated.	1153
ALR4090W Active zone zone name in fabric fabric name or WWN is missing.	1153
ALR4091I Active zone zone name in fabric fabric name or WWN has been discovered.	1153
ALR4092I Active zoneset zoneset name in fabric fabric name or WWN has been discovered.	1153
ALR4093E ZoneSet zoneset name in fabric fabric name or WWN has been deactivated. ZoneSet zoneset name has been activated.	1154
ALR4094I Active zone zone name in fabric fabric name or WWN has been rediscovered.	1154
ALR4095I Active zoneset zoneset name in fabric fabric name or WWN has been rediscovered.	1154
ALR4096I Zone member zone member name has been added to active zone zone name in fabric fabric name or WWN .	1154
ALR4097I Zone member zone member name has been removed from active zone zone name in fabric fabric name or WWN .	1154
ALR4098I Zone member zone member name has been readded to active zone zone name in fabric fabric name or WWN .	1155
ALR4099I Zone zone name has been added to active zone set zone set name in fabric fabric name or WWN .	1155
ALR4103W The performance monitor's primary process is experiencing a high memory utilization.	1155
ALR4104W A database used by the system is reporting an alarm: value.	1155
ALR4105W Device server configuration should be changed to improve performance: value.	1155
ALR4106W The IBM Spectrum Control server is receiving a high number of external type of events received, which is either CIM for CIM indications or SNMP for SNMP traps. events, which may cause temporary performance degradation.	1156
ALR4107I Zone zone name has been removed from active zone set zone set name in fabric fabric name or WWN .	1156
ALR4108I The server server name at host host name successfully connected to the database after previous attempts failed.	1156
ALR4112E The server server name at host host name failed to connect to the database.	1156
ALR4109I Alias zone alias has been added to active zone zone name in fabric fabric name or WWN .	1157
ALR4113E The IBM Spectrum Control server failed to connect to the database.	1157
ALR4110W Alias alias namehas been removed from active zone active zone name in fabric fabric name .	1157
ALR4111I Alias alias namehas been readded to active zone active zone name in fabric fabric name .	1157
ALR4197W A new connection is detected.	1157
ALR4198W The state for connection from initial state to final state has changed.	1158
ALR4199W The state for Fabric fabric name has changed.	1158
ALR4224W The state has changed for Node node name .	1158
ALR4225I Node node name has been discovered.	1158
ALR4226W Node node name has gone offline.	1158
ALR4227I Node node name has gone online.	1158
ALR4241E Subsystem storage subsystem name has gone offline.	1159
ALR4242I Subsystem storage subsystem name has gone online.	1159
ALR4243W The subsystem version has changed from initial version to new version on Subsystem storage subsystem name .	1159

ALR4244W The allocated capacity has changed from initial capacity to new capacity on Subsystem storage subsystem name .	1159
ALR4245W The available capacity has changed from initial capacity to new capacity on Subsystem storage subsystem name .	1159
ALR4246W Back-end capacity has changed from initial capacity to new capacity on Subsystem storage subsystem name .	1160
ALR4247W Back-end controller back-end controller name for owning storage subsystem name has gone offline.	1160
ALR4248I Back-end controller back-end controller name for owning storage subsystem name has gone online.	1160
ALR4249W Volume volume name on storage subsystem name has gone offline.	1160
ALR4250I Volume volume name on storage subsystem name has gone online.	1160
ALR4251W The capacity has changed from initial capacity to new capacity for Volume volume name on Subsystem storage subsystem name .	1161
ALR4252W The state for Pool pool name on Subsystem storage subsystem name has changed to not detectable .	1161
ALR4253I Pool pool name on Subsystem storage subsystem name has been discovered.	1161
ALR4254W Pool pool name on Subsystem storage subsystem name has gone offline.	1161
ALR4255I Pool pool name on Subsystem storage subsystem name has gone online.	1161
ALR4256W The pool capacity has changed from initial capacity to new capacity for Pool pool name on Subsystem storage subsystem name .	1162
ALR4257W The pool available space has changed from initial capacity to new capacity for Pool pool name on Subsystem storage subsystem name .	1162
ALR4273E Server server name has gone offline.	1162
ALR4274I Server server name has gone online.	1162
ALR4278W The property for Subsystem storage subsystem name has changed.	1162
ALR4300W The use count for Disk Drive disk drive name on Subsystem subsystem name has changed from initial use count to final use count.	1163
ALR4301W Disk Drive disk drive name on Subsystem subsystem name has gone offline.	1163
ALR4304W The state for Back-end Controller back-end controller for subsystem has changed from initial state to final state .	1163
ALR4305W The WWPN path count for Back-end Controller back-end controller for system name has changed from initial state to final state .	1163
ALR4312W Notification has received from external device device name	1163
ALR4313W The endpoint version has changed from initial state to final state on Endpoint	1164
ALR4314W Entity entity name has been discovered.	1164
ALR4315I Hypervisor hypervisor name has been discovered.	1164
ALR4316I Virtual machine vm name was added to hypervisor hypervisor name .	1164
ALR4317W Virtual machine vm name was removed from hypervisor hypervisor name .	1164
ALR4318W New unmanaged hypervisor discovered.	1164
ALR4319W Virtual machine added.	1165
ALR4320W Virtual machine removed.	1165
ALR4321W Hypervisor hypervisor missing.	1165
ALR4322W Hypervisor missing.	1165
ALR4323I New disk disk name discovered for system system name .	1165
ALR4324W Disk disk not found for system system .	1166
ALR4325W New volume volume name discovered for system system name .	1166
ALR4326W Volume volume not found for system system.	1166
ALR4327W Zone Alias to Member Change	1166
ALR4328W The association between Zone Alias zone alias and Member member has changed.	1166
ALR4329I Zone member zone member has been added to zone alias zone alias in fabric fabric .	1167
ALR4330W Zone member zone member has been removed from zone alias zone alias in fabric fabric .	1167
ALR4331I Zone member zone member has been read added to zone alias zone alias in fabric fabric .	1167
ALR4332W Storage Resource Agent Deployment Failed	1167
ALR4333W Replication Session State Change alert received.	1167
ALR4334W Replication Configuration Change alert received.	1168
ALR4335W Replication Suspending Event Notification alert received.	1168
ALR4336W Replication Communication Failure alert received.	1168
ALR4337W Replication Management Server State Change alert received	1168
ALR4338W Replication PPRC Path State Change alert received.	1168
ALR4339W The IBM Spectrum Control for Replication resource resource name triggered an alert with the following message: message text	1168
ALR4353W number of affected datapaths Data Paths from Host host name to Volume volume name on Subsystem subsystem name are no longer available.	1169
ALR4354I number of affected datapaths Data Paths from Host host name to Volume volume name on Subsystem subsystem name have been discovered..	1169
ALR4356E The mount state of specified file system changed to error level.	1169
ALR4358I The mount state of specified file system changed to normal level.	1169
ALR4359E The CPU usage reached the error level.	1170
ALR4360W The CPU usage reached the warning level.	1170
ALR4361I The CPU usage reached the normal level.	1170
ALR4362E The memory usage reached the error level.	1170
ALR4364I The memory usage reached the normal level.	1170
ALR4365I The clustered CIFS is active.	1170
ALR4366W The clustered CIFS is disabled.	1171
ALR4367E The clustered CIFS status reached the error level.	1171
ALR4368I The IBM Spectrum Scale is active.	1171
ALR4369W The IBM Spectrum Scale status reached the warning level.	1171
ALR4370E The IBM Spectrum Scale is down.	1171
ALR4427I The file system file system has been detected on device type device device display name.	1172
ALR4429I The capacity of file system file system has changed from previous capacity to current capacity on device type device NAS display name.	1172
ALR4433W The free space on file system file system has fallen below the threshold value of threshold on device type device nas display name. The free space is current free space value or relative free space value of the file system capacity.	1172
ALR4440W The state of the device type node type node node name changed from old state to new state on device display name.	1172
ALR4441I The state of the device type node type node node name changed from old state to new state on device display name.	1172
ALR4442W The IBM Spectrum Scale status of device type node node name changed from old state to new state on device display name.	1173
ALR4443I The IBM Spectrum Scale status of device type node node name changed from old state to new state on device display name.	1173

ALR4447I The fileset fileset was detected on device type device device name.	1173
ALR4448I The fileset fileset was linked to path for device type device device name.	1173
ALR4455I The fileset fileset was unlinked on device type device device name.	1173
ALR4458W The number of free inodes on file system path has fallen below the threshold value of threshold on device type device device name. There are current value free inodes or current value relative to maximum of the maximum inodes.	1174
ALR4460I Export export name detected on device type device device name with path path.	1174
ALR4461W The state of export export name changed from previous state to current state on device type device device name.	1174
ALR4462W Export export name was reconfigured on device type device device name. Path changed from previous path to current path.	1174
ALR4463W Export export name was reconfigured on device type device device name. Protocols changed from previous list of protocols to current list of protocols.	1174
ALR4470W Export export name is missing from device type device device name..	1175
ALR4471I Export export name was rediscovered on device type device device name..	1175
ALR4474W Fileset fileset name is missing from device type device device name..	1175
ALR4475I Fileset fileset name was rediscovered on device type device device name..	1175
ALR4478W File system File system name is missing from device type device device name..	1175
ALR4479I File system File system name was rediscovered on device type device device name..	1176
ALR4482W A Quota type Quota limit type quota was violated for the path file system on the device device name device type system. Quota type resource name is consuming usage and the Quota limit type limit is threshold.	1176
ALR4496I New quota detected on file system path of Device type device Device name.	1176
ALR4385E The status of NSD NSD name reached error level.	1176
ALR4386W The status of NSD NSD name reached warning level.	1176
ALR4387I The status of NSD NSD name was set back to normal level.	1177
ALR4503I New NSD NSD name has been detected on device type device device display name.	1177
ALR4505W NSD NSD name is missing from device type device device name.	1177
ALR4507W The state of NSD NSD name changed from previous state to current state on device type device device name.	1177
ALR4511E Alert condition for nodes has been selected. Select only nodes.	1177
ALR4512E Alert condition for clusters has been selected. Select only clusters.	1178
ALR4513E Alert condition for NSD has been selected. Select only NSD.	1178
ALR4514E Alert condition for File set has been selected. Select only File set.	1178
ALR1022M A new unmanaged server or cluster discovered.	1178
ALR1294W The server or cluster has gone offline.	1178
ALR1295W The server or cluster has gone online.	1178
ALR1296W The server or cluster property has changed.	1179
ALR1245W A node state has changed.	1179
ALR1246W A node was discovered.	1179
ALR4528I Cluster was discovered.	1179
ALR4529I Cluster was removed.	1179
ALR4530I Cluster was rediscovered.	1180
ALR0078W =Performance monitor for device value failed to collect new data using data source value.	1180
ALR4391I Node node name is selected as cache gateway node.	1180
ALR4392I Node node name is unselected as cache gateway node.	1180
ALR4393I Home system home system name detected on device type device device name with path path.	1180
ALR4394I Home system has been removed from fileset fileset name on device type device device name.	1180
ALR4395W Home system home system name is missing from device type device device name.	1181
ALR4396I Home system home system name was rediscovered on device type device device name..	1181
ALR4397I Cache fileset cache fileset name detected on device type device device name.	1181
ALR4398W Cache fileset cache fileset name is missing from device type device device name.	1181
ALR4399I Cache fileset cache fileset name was rediscovered on device type device device name.	1181
ALR4400I Cache fileset name has changed from cache fileset name to cache fileset name on device type device device name.	1182
ALR4401I Cache fileset cache fileset name state has changed from old value to new value on device type device device name.	1182
ALR4402I Cache fileset cache fileset name mode has changed from old value to new value on device type device device name.	1182
ALR4403I Cache client cluster name is added to home system home system name on device type device device name.	1182
ALR4404I Cache client cluster name has been removed from home system home system name on device type device device name.	1182
ALR4541E The available space is too low for pool pool name on storage system storage system name. The measured value pool available space violates the critical boundary of user defined threshold value.	1183
ALR4542W The available space is too low for pool pool name on storage system storage system name. The measured value pool available space violates the warning boundary of user defined threshold value.	1183
ALR4543E The allocation is too high for pool pool name on storage system storage system name. The measured value pool virtual allocation violates the critical boundary of user defined threshold value.	1183
ALR4544W The allocation is too high for pool pool name on storage system storage system name. The measured value pool virtual allocation violates the warning boundary of user defined threshold value.	1183
ALR4545E The shortfall percentage is too high for pool pool name on storage system storage system name. The measured value pool shortfall percentage violates the critical boundary of user defined threshold value.	1184
ALR4546W The shortfall percentage is too high for pool pool name on storage system storage system name. The measured value pool shortfall percentage violates the warning boundary of user defined threshold value.	1184
ALR4547I VMWare Cluster cluster name discovered on hypervisor hypervisor name.	1184
ALR4548W VMWare Cluster cluster name removed from hypervisor hypervisor name.	1184
ALR4549I New cluster hypervisor relationship discovered.	1184
ALR4550W Cluster hypervisor relationship removed.	1185
ALR4551I Hypervisor hypervisor name was moved from VMWare Cluster old cluster nameto VMWare Cluster new cluster name.	1185
ALR4552I Cluster hypervisor relationship moved.	1185
ALR4600I Fabric Name fabric WWN changed to fabric WWN .	1185
ALR1349I A new path path name was discovered for disk disk name on host host name.	1185

ALR1350W The path path name was not found for disk disk name on host host name.	1186
ALR1351W The path path name for disk disk name on host host name is disconnected.	1186
ALR4604I The home system home system was linked to path on resource type resource resource name.	1186
ALR4605I The home system home system was unlinked from a path on resource type resource resource name.	1186
ALR1352E The status of disk disk name on server server name has degraded to Error from former status.	1186
ALR1353W The status of disk disk name on server server name has degraded to Warning from former status.	1187
ALR1354I The status of disk disk name on server server name has improved to Normal from former status.	1187
ALR4625I New entity type, entity name, added to system type system name .	1187
BPCCA - Data collector installation messages	1187
BPCCA0001I The data collector started, connected to the storage management service, and is ready to process requests from the storage management service.	1188
BPCCA0002E The data collector failed to connect to the storage management service at server_url.	1188
BPCCA0003E The data collector started but detected a problem with the directory directory_name and must stop.	1188
BPCCA0004E The data collector cannot run because it is not configured correctly.	1188
BPCCA0005E The data collector failed to connect to the storage management service at server_url because the host name could not be resolved.	1189
BPCCA0006E The data collector failed to connect to the storage management service at server_url because of an unknown error.	1189
BPCCA0007E The data collector failed to connect to a service from the storage management system.	1189
BPCCA0008E The data collector failed to connect to the storage management service because of invalid credentials.	1189
BPCCA0009I The data collector connected to the storage management service. The data collector had failed to connect since date_and_time.	1190
BPCCA0010E The data collector in the directory_name directory of the host host_name was running and an attempt was made to start a second instance of the same data collector. The second instance of the data collector stopped.	1190
BPCCA0011I The data collector stopped because a user requested it to shut down.	1190
BPCCA0012I The data collector stopped to enable the installation of an upgraded version of the data collector.	1190
BPCCA0013E The storage management service did not allow the data collector to connect because another data collector was already connected to the service.	1190
BPCCA0100I The updateCollector utility started.	1191
BPCCA0101E The collector directory was not specified in the collectorDirectory.properties file.	1191
BPCCA0102E The collector directory directory_name that was specified in the collectorDirectory.properties file is invalid. The collector directory is the directory to which the updateCollector utility must copy the upgrade image files.	1191
BPCCA0103E The collector directory directory_name that was specified in the collectorDirectory.properties file cannot be used as the collector directory.	1191
BPCCA0104E The updateCollector utility started but there was a problem with the upgrade image directory current_directory.	1191
BPCCA0105E The collector directory directory_name that was specified in the collectorDirectory.properties file is a subdirectory of the upgrade image directory upgrade_image_directory_name.	1192
BPCCA0106E Cannot upgrade the data collector in the collector_directory directory because the directory contains the following locked files: locked_files_list	1192
BPCCA0107I The content of the directory collector_directory will be deleted and replaced with subdirectories and files from the upgrade_image_directory directory. Some configuration files, the log directory, and the contents of the log directory will not be deleted.	1192
BPCCA0108E The data collector service cannot be uninstalled from the operating system. The upgrade process cannot be completed.	1192
BPCCA0109E The updateCollector utility could not upgrade the data collector. The data collector service is now in an inconsistent state.	1193
BPCCA0110E The contents of the collector_directory directory could not be deleted. The upgrade process cannot be completed. The data collector might be in an inconsistent state.	1193
BPCCA0111E The files and directories of the data collector from the directory upgrade_image_directory could not be copied into the directory collector_directory. The upgrade process cannot be continued. The data collector might be in an inconsistent state.	1193
BPCCA0112I The data collector in the directory collector_directory was upgraded successfully to version downloaded_version.	1193
BPCCA0113E The data collector in the directory collector_directory could not be upgraded.	1194
BPCCA0114I The data collector was upgraded to the new version and will start automatically.	1194
BPCCA0115I The attempt to upgrade the data collector failed. The existing data collector will start automatically.	1194
BPCCA0116E The attempt to upgrade the data collector failed. You must download and install the latest version of the data collector.	1194
BPCCA0117E The upgraded data collector could not be installed as a service on the operating system. You must install the new data collector service manually.	1194
BPCCA0118I The data collector was upgraded to the new version and started successfully.	1195
BPCCA0119I The data collector could not be upgraded, but was not modified. The existing data collector was restarted successfully.	1195
BPCCA0120E The upgraded data collector did not start.	1195
BPCCA0121E The existing data collector did not restart.	1195
BPCCA0122E The data collector cannot authenticate to the HTTPS proxy server proxy_server_hostname.	1195
BPCIN - Spectrum Control installation messages	1196
BPCIN0001I The system is installing IBM Spectrum Control.	1201
BPCIN0002I The system completed the installation of IBM Spectrum Control.	1201
BPCIN0003E The installation program could not find the file file_name in the installation image. For more information, go to the IBM Knowledge Center and search on the message code.	1201
BPCIN0004E An error occurred during the installation of the component_name component. Review the log files in the log_file_directory_name directory for additional information.	1201
BPCIN0005E Invalid characters "characters" were found in the installation path "path".	1202
BPCIN0006E The location "location" that was specified for the installation is not empty. It might contain hidden items.	1202
BPCIN0007E Directory directory is not writable.	1202
BPCIN0008E The installation location specified is blank. Enter an installation location.	1203
BPCIN0009E The host name or IP address is not valid for name_with_spaces.	1203
BPCIN0010E The port number field is blank. Enter a valid port number.	1203
BPCIN0011E Enter a fully qualified host name or IP address.	1203
BPCIN0012E Enter a valid port number.	1203
BPCIN0013E Enter a port number from 1 to 65535.	1203
BPCIN0014E The user name cannot be blank or contain spaces. Enter a valid user name.	1204
BPCIN0015E The location for the license key file is blank. Enter a file location.	1204
BPCIN0016E The file that was specified does not exist. Enter a valid license key file.	1204
BPCIN0017E The file that was specified is not a valid license key file.	1204
BPCIN0018E The installation path cannot end with the characters "endCharacter".	1204
BPCIN0019I The system is installing the component component.	1205

BPCIN0020I The system completed the installation of the component component.	1205
BPCIN0021E An unexpected error occurred during the prevalidation of the component.	1205
BPCIN0022E The port range port_start - port_end is not available because the port or ports port are already used by other applications. The next available port range is available_start_port - available_end_port.	1205
BPCIN0023E The password is incorrect. It cannot be blank or contain spaces. Enter a valid password.	1206
BPCIN0024E The user name userID cannot contain the following special characters: characters.	1206
BPCIN0025E The password password cannot contain spaces or any of the following special characters: characters.	1206
BPCIN0026E The installation location selected does not have enough space.	1206
BPCIN0027E The user name userID is not in the administrative group adminGroup.	1206
BPCIN0028E The validation for user name userID has failed. Check to see if this user name exists.	1207
BPCIN0029E The password that was entered does not match the password for user user_name on the system. Please try again.	1207
BPCIN0031E The configuration of component failed. Review the installation log file.	1207
BPCIN0032E The host name field is empty. Enter the fully qualified host name or IP address.	1207
BPCIN0033E The specified host name or IP address cannot be identified.	1207
BPCIN0034E The installation location is not an absolute path. Enter an absolute path.	1208
BPCIN0035W The installation program cannot validate the host name because the fully qualified domain name (FQDN) cannot be retrieved for the host.	1208
BPCIN0043E The last port number must be a port number from 1 to 65535.	1208
BPCIN0044E The installation location that was specified uses a Windows reserved name. Enter a different installation location.	1208
BPCIN0045E The installation location cannot contain special shell characters special_characters in the installation path path.	1209
BPCIN0046E The installation location contains special characters special_characters that are not supported by the operating system in the installation path path.	1209
BPCIN0047E The IBM Db2 database manager must be active to continue the installation process.	1209
BPCIN0048E IBM Spectrum Control cannot be installed because Db2 is not installed on the system or the Db2 configuration is not valid.	1209
BPCIN0049E The Db2 db2Version is not supported. The minimum supported Db2 versions are minDB2Version.	1209
BPCIN0050E An error occurred when the installation program tried to verify that the Db2 database manager is running. The error message is: error_message.	1210
BPCIN0051E An error occurred when the installation program tried to find an available database named dbName.	1210
BPCIN0052I The name of the database to be created is dbName.	1210
BPCIN0053E An error occurred when creating the database dbName. Review the log files for more information.	1210
BPCIN0055E Db2 is not installed or the Db2 profile was not sourced before installing IBM Spectrum Control.	1211
BPCIN0056E The user name userID does not have write permission on the default database path configuration parameter DFTDBPATH: dftdbpath.	1211
BPCIN0057E The user name userID is not in an operating system group that has Db2 SYSADM authority. Before you run your DB2 installation, validate that the Windows Server Hostname is 15 characters or less in length.	1211
BPCIN0059W There are long file names in the IBM Spectrum Control installation images.	1211
BPCIN0060E An error occurred during the uninstallation of component. Review the log files for more information.	1212
BPCIN0061E An invalid host name or IP address was specified for the Data server.	1212
BPCIN0062E The IPv6 internet protocol is not enabled on the specified host computer.	1212
BPCIN0063E The Data server is not running. For more information, go to the IBM Knowledge Center and search on the message code.	1212
BPCIN0064E The Data Server is not running at the specified host address or port: host host_address, port port. For more information, go to the IBM Knowledge Center and search for the message code.	1213
BPCIN0066E Errors occurred during the installation of the IBM Spectrum Control GUI. Review the log files for more information.	1213
BPCIN0068E Errors occurred during the uninstallation of the IBM Spectrum Control GUI. Review the log files for more information.	1213
BPCIN0069E Errors occurred during the configuration of Tivoli Common Reporting for IBM Spectrum Control. Review the log files for more information.	1214
BPCIN0070E Errors occurred during the configuration of the IBM Spectrum Control data model in Tivoli Common Reporting. Review the log files for more information.	1214
BPCIN0071E An error occurred because the installation program could not find the Db2 DFTDBPATH variable. For more information, go to the IBM Knowledge Center and search for the message code.	1215
BPCIN0072E The validation for user name userID has failed. Check to see if this user name exists or if Db2 is running.	1215
BPCIN0074E The port range validation failed because the port value is not numeric.	1215
BPCIN0075E An invalid GUID 0xFFFFFFFFFFFFFFFFFFFFFFFF was found. Update or uninstall the GUID.	1215
BPCIN0076E IBM Spectrum Control could not read the GUID. See the GUID installation log for an explanation of the error.	1215
BPCIN0077E The name specified for the database is not valid because it contains a space or is blank.	1216
BPCIN0078E The database name database_name is not valid. The name can only contain the following characters: a-z, A-Z or 0-9.	1216
BPCIN0079E The first character or characters in the database name are not valid. The name must not begin with a number or the letters: SYS, DBM, or IBM.	1216
BPCIN0080E The database name specified is not valid because an existing database has the same name: database_name.	1217
BPCIN0081E The database name database_name is too long. The name can be 1 - 8 characters in length with no spaces.	1217
BPCIN0082E The specified host name or IP address is not a remote host. Enter a remote host name or IP address.	1217
BPCIN0083E The server could not connect to the remote database. Verify that the port is correct and that the database is running on the remote server. Also verify that the internet protocol connection between the server and remote database is compatible.	1217
BPCIN0084E The database name is not valid because there is no database on the specified server with this name.	1217
BPCIN0085E The installation program was unable to connect to the remote database. Check the remote server to verify that you can connect to Db2. For more information about Db2 connection issues, go to the IBM Knowledge Center at http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome .	1218
BPCIN0086E The connection to the remote database failed because the user name or password is not valid.	1218
BPCIN0087E The remote host name or IP address is not valid.	1218
BPCIN0088E The database repository was not found on the remote server.	1218
BPCIN0089E The version "retrivedVersion" of the remote database repository is not at the correct level. Upgrade the remote database repository to version "requiredVersion", then upgrade the remaining components.	1219
BPCIN0090E The name specified for the database is not valid because it contains a space.	1219
BPCIN0091E The file path file_path specified is not an absolute path or is not a directory or the partition does not exists.	1219
BPCIN0092E The first directory directory_name specified is not valid.	1219
BPCIN0093E The path string path_string is too long. The path string cannot be longer than 242 bytes.	1219
BPCIN0094E The user name user_name does not have write permission on the specified database path: database_path.	1220
BPCIN0095E The user name user_name does not have write permission on the specified log location: log_location.	1220
BPCIN0096E The database path cannot be blank. Enter a valid database path.	1220

BPCIN0097E The log location cannot be blank. Enter a valid log location.	1220
BPCIN0098E Enter 10 or fewer paths.	1220
BPCIN0099E You have entered pathNumber directories. A maximum of 10 directories can be specified.	1221
BPCIN0100W The database log files and database are in the same location. Click Yes to change the database log files or database location. Click No to ignore this message.	1221
BPCIN0101E The location "location" that was specified contains Db2 logs or log files.	1221
BPCIN0102E The Db2 profile for the "db2_instance" instance was not run before installing IBM Spectrum Control.	1221
BPCIN0103E The user name userID is not in the system group adminGroup.	1221
BPCIN0104I Resuming a failed install and installing the remaining components.	1222
BPCIN0105E A reboot must be done before continuing with the installation.	1222
BPCIN0106E There are invalid header files in the installation images. You may might be using the AIX tar program instead of the GNU tar program to extract files from the installation images. For more information, see IBM Knowledge Center and search by the error message code.	1222
BPCIN0107E The database path exists in the database path list. You must specify a unique database path.	1222
BPCIN0108E An error occurred during the upgrade of the component. Review the log files in the following directory for an explanation of the error: location.	1223
BPCIN0109E An unexpected error occurred. IBM Spectrum Control cannot resolve this error. For more information, review the log files and go to the IBM Knowledge Center.	1223
BPCIN0112I Upgrading the components:upgrade_components.	1223
BPCIN0116I The system is upgrading the component component.	1224
BPCIN0117I The system completed the upgrade of the component component.	1224
BPCIN0120E The validation of user name and password could not be completed	1224
BPCIN0121E The installation type entered is not supported.	1224
BPCIN0122E The TPCCCommon.dll or libTPCCCommon.so library could not be found or loaded.	1224
BPCIN0123E The user has no administrative rights to install or uninstall IBM Spectrum Control.	1225
BPCIN0124E The operating system cannot find the Db2 Service name.	1225
BPCIN0125E The installation program cannot find the Db2 instance name. Check the Db2 instance configuration.	1225
BPCIN0126E Unable to find Db2 installation path. Check your Db2 configuration.	1225
BPCIN0127E The service name configuration parameter cannot be queried from the Db2 Database Manager.	1225
BPCIN0129E The version that is installed cannot be upgraded because it does not meet the minimum build version for the database repository. The minimum build version supported is build_version.	1226
BPCIN0130E The database administrator password in IBM Spectrum Control does not match the database administrator password for Db2. Run the changepasswords tool to update the database administrator password.	1226
BPCIN0131E An invalid host name or IP address was specified for the Device Server.	1226
BPCIN0132E The Device Server is not running at the specified host address or port: host host_address, port port.	1227
BPCIN0134E The installation program could not stop the . Please manually stop the .	1227
BPCIN0135E Tivoli Storage Productivity Center version is installed. Before you can upgrade to version 5, you must upgrade Tivoli Storage Productivity Center to version 4.	1227
BPCIN0136E The version "retrivedVersion" of the remote database repository is not at the correct build level. Upgrade the remote database repository to the correct build level, then upgrade the remaining version "requiredVersion" components.	1227
BPCIN0137E The version "retrivedVersion" of the remote database repository is not at the correct build level. Ensure that the installation image for the remaining version "requiredVersion" components is at the same build level as the remote database repository, then upgrade the remaining components.	1228
BPCIN0138E The common user name password does not match the password for the existing user name user_name on the system. Run the changepasswords tool to change the common user name password for IBM Spectrum Control.	1228
BPCIN0139I The User Migration Tool completed successfully.	1228
BPCIN0140E The User Migration Tool could not be started. Start the User Migration Tool from the graphical user interface after the upgrade is finished.	1228
BPCIN0143E Errors occurred during the upgrade of the GUI. Review the log files for more information.	1229
BPCIN0146E The installation language specified is not a supported language.	1229
BPCIN0148E IBM Spectrum Control cannot be installed, because the physical memory size on this computer is too small. The minimum memory size is minMemoryProduction. Increase the physical memory on this computer and run the installation program again.	1229
BPCIN0149W The physical memory size on this computer is below the minimum requirements that are specified for a production system. The minimum memory size for a production system is minMemoryProduction. You must increase the physical memory on this computer and run the installation program again. If you install IBM Spectrum Control with lower memory, you can only use it in an evaluation environment.	1230
BPCIN0150E The installation program cannot determine the amount of physical memory on the system.	1230
BPCIN0151E You tried to install IBM Spectrum Control on an unsupported operating system. For more information on the supported operating systems, go to following link.	1230
BPCIN0152E IBM Spectrum Control cannot determine if the operating system is supported.	1230
BPCIN0153E The installation program could not rename the jre folder. To continue the installation, stop all Java processes that access the jre directory: jre_folder	1231
BPCIN0159E The web server data source creation failed.	1231
BPCIN0163E The service name configuration parameter cannot be queried from the Db2 Database Manager. Reboot the machine after Db2 installation. SQL Error Message is sqlerrmsg	1231
BPCIN0164E The service name configuration parameter cannot be queried from the Db2 Database Manager. SQL Error Message is sqlerrmsg	1232
BPCIN0165E The password for user name has expired. You must change the password, or select another user name, to install IBM Spectrum Control.	1232
BPCIN0166W The password for user name cannot be checked for expiration. Please ensure that it is not expired.	1232
BPCIN0167E You cannot upgrade the license to the same or lower level.	1232
BPCIN0168E An error occurred while checking the user name and password. Review the log files for more information.	1232
BPCIN0169E The location "location" that should be specified for the license is not correct.	1233
BPCIN0170E IBM Spectrum Control cannot upgrade the license because the upgrade must be done on the server system.	1233
BPCIN0173E An error occurred during the deployment of the file_name file.	1233
BPCIN0176E An error occurred during Db2 catalog creation, and the catalog was not created.	1234
BPCIN0177E You cannot upgrade IBM Spectrum Control with a license key file that is at a lower level than the installed license key file.	1234
BPCIN0178E The provided license key file_name is invalid.	1234
BPCIN0179E An error occurred during Db2 catalog deletion, and the catalog was not removed.	1234
BPCIN0181E The web server data source testing has failed.	1235
BPCIN0182E TPC-GUI.war data source testing failed.	1235

BPCIN0184E The remote database repository already contains data. Install a new IBM Spectrum Control remote database repository first, then install the remaining IBM Spectrum Control components.	1235
BPCIN0185E The remote database repository for version oldVersion is incompatible with IBM Spectrum Control Version requiredVersion. Upgrade remote database repository to Version requiredVersion first, and then install the remaining IBM Spectrum Control components.	1235
BPCIN0190E You cannot use non-standard characters, such as a space between characters or an underscore, in a host name. You must enter a host name with standard characters and try again.	1235
BPCIN0191E You cannot install Db2 in a directory that starts with the letter a.	1236
BPCIN0193E The version "retrivedVersion" of the remote database repository is not at the correct level. Install a new Version "requiredVersion" remote database repository at the correct level, then install the remaining IBM Spectrum Control components.	1236
BPCIN0195E The version "retrivedVersion" of the database repository dbHost is not at the correct build level. Install a new Version "requiredVersion" database repository at the correct build level, then install the remaining IBM Spectrum Control components.	1236
BPCIN0198E The path to installation image "installDirectory" contains the invalid character "unallowedCharater". You must change the directory name so that it contains valid characters.	1236
BPCIN0199E The installation program cannot validate the host name because the fully qualified domain name (FQDN) cannot be retrieved for the host.	1237
BPCIN0200E The installation program does not allow host name specified as IP address. Please ensure that a fully qualified domain name (FQDN) is provided.	1237
BPCIN0202E Runtime errors have occurred during the IBM Spectrum Control preinstallation process. The installation program cannot recover from this error.	1237
BPCIN0203E The installation program does not support backslash character in the user name. In case it is a Windows domain account, please specify just the user name without using the "Domain_name\\" prefix.	1237
BPCIN0205E The IBM Spectrum Control installation program could not find the directory "installDirectory" on the installation image. For more information, go to the IBM Knowledge Center and search on the message code.	1238
BPCIN0206E The IBM Spectrum Control installation program was unable to retrieve a fully qualified domain name (FQDN) for the host. You must configure the host system with an FQDN.	1238
BPCIN0207E You cannot upgrade when the stand-alone GUI is running. Stop the stand-alone GUI and click OK to continue the upgrade or click Quit to exit the installation program.	1238
BPCIN0208E The fully qualified domain name(FQDN) retrieved for the host contains non-standard characters, such as a space between characters or an underscore. You must configure the host system with an FQDN that contains standard characters.	1238
BPCIN0209E The installation program does not allow a user that is present in both the Windows Domain and the local Operating System repositories. Install with a user name that is only present in one of these repositories.	1239
BPCIN0210E The Db2 "sourcedb2profile" profile has not been loaded in the .profile file for user commonUser. This profile must be loaded before you install IBM Spectrum Control.	1239
BPCIN0211E The technology level or the service pack level of this operating system is not supported and must be upgraded to a supported level. The detected operating system version is os_version	1239
BPCIN0214E Cannot complete a fresh install without a valid license.	1239
BPCIN0215E The path extractorDirectory where IBM Spectrum Control is extracted has extractorDirLength characters and the maximum number of characters allowed is 260. You must shorten this path.	1240
BPCIN0217E The Windows registry check indicates that .NET 3.5 or higher is not available on this Windows server. You must install .NET 3.5 or higher to continue. To install .NET 3.5 type the following commands in a 64 bit windows powershell: Import-Module ServerManager Add-WindowsFeature as-net-framework	1240
BPCIN0219E The domain configuration is invalid. To resolve the issue, complete the following steps: Disable the Windows Firewall service. Start or restart the Computer Browser service on this domain member computer and on the domain controller computer. If the service has a Stopped or Disabled status on the domain controller computer, you must restart Computer Browser service on the domain member computer after you start the service on the domain controller computer. In a Windows command window, run the net view command and verify that there are no errors. Reinstall IBM Spectrum Control.	1240
BPCIN0220E The current login user loginuser is not an administrator or a member of the domain administrator group. The installation program cannot start.	1240
BPCIN0221E The user name userID is not a part of the local administrator group.	1241
BPCIN0222E The current login user loginuser is not an administrator or a member of the local administrator group. The installation program cannot continue.	1241
BPCIN0223E The current login user loginuser is not a member of the local Db2 administrator group. The installation program cannot start.	1241
BPCIN0224E The user userID is a domain account and cannot be used to log in to Db2. You must enter a separate user name for Db2.	1241
BPCIN0225E The current logged in user loginuser does not have Db2 SYSADM authority. To provide the user with Db2 SYSADM authority, log in by using a user name with SYSADM authority and run the following commands: db2cmd db2set -g DB2_GRP_LOOKUP=local,TOKENLOCAL db2 force application all db2stop db2start	1242
BPCIN0226E Login as a windows domain user in order to install IBM Spectrum Control using a windows domain account.	1242
BPCIN0229E The Monitoring Agent for Windows OS - Watchdog and Monitoring Agent for Windows OS - Primary services must be stopped before you can continue. After you upgrade, you must restart these services.	1242
BPCIN0230E The Monitoring Agent for Windows OS - Watchdog and Monitoring Agent for Windows OS - Primary services must be stopped before you can continue. After you uninstall IBM Spectrum Control, you must restart these services if the required reboot is postponed.	1242
BPCIN0231E Tivoli Common Reporting is not installed on your system. You must install it to continue with the IBM Spectrum Control installation.	1243
BPCIN0233E Jazz for Service Management is not installed in the specified location. Reenter the correct installation location for Jazz for Service Management.	1243
BPCIN0234E The Jazz for Service Management user credentials that you entered were incorrect. Check the user credentials and try again.	1243
BPCIN0235W The Jazz for Service Management and Tivoli Common Reporting servers are not running. To start the servers, run the following command: file It takes a few minutes for these servers to start up and initialize. Wait for a few minutes before resuming the installation.	1243
BPCIN0236W The Tivoli Common Reporting server is not running. If you have already started the server , wait for a few more minutes. It takes a while for the server to start up and initialize. Otherwise, you can start the server by running this command: file	1244
BPCIN0238E The Tivoli Common Reporting server at tcrllocation cannot be reached. If the Jazz for Service Management server has not completed startup, wait for a few more minutes and then click OK. It takes a while for the server to start and initialize. Otherwise, if the Jazz for Service Management server is started but Tivoli Common Reporting still cannot be reached, restart the Jazz for Service Management server using the following commands: stopServer -username username -password password startServer	1244
BPCIN0239E The Tivoli Common Reporting server configuration cannot be exported.	1244
BPCIN0240E The Tivoli Common Reporting server configuration file "configuration file" was not created.	1244
BPCIN0241E The Tivoli Common Reporting configuration cannot be upgraded.	1245
BPCIN0242E The Windows registry check indicates that .NET 3.5 is not available on this Windows 2012 server. You must install .NET 3.5 to continue. To install .NET 3.5 type the following commands in a 64 bit windows powershell: Import-Module ServerManager Add-WindowsFeature as-net-framework	1245
BPCIN0244E An error occurred while enumerating the local administrator group membership. On the current computer, on the Properties page of this group, remove the user names that are displayed with a security ID (SID). An example of a SID is S-1-5-21-337177553-1671989427-887411491-500 and is used instead of a user name.	1245
BPCIN0245E An error occurred while enumerating the local Db2 administrator group membership. On the current computer, on the Properties page of this group, remove the user names that are displayed with a security ID (SID). An example of a SID is S-1-5-21-337177553-1671989427-887411491-500 and is used	1245

instead of a user name.	1245
BPCIN0246E Tivoli Storage Productivity Center version is installed. Before you can upgrade to version 5, you must upgrade Tivoli Storage Productivity Center to version 4.	1246
BPCIN0247E An error occurred during the domain check prevalidation process. Verify that the domain controller computer is available and then restart this domain member machine.	1246
BPCIN0248E The IBM Spectrum Control installation program supports only fully qualified user names on Windows domain member machines. Specify the user name userID by using the "Domain_name\\" or the "Machine_name\\" prefix. The detected domain name is "domainName".	1246
BPCIN0250W Before you can upgrade the database repository, you must first stop the Data server, Device server, Alert server, Export server, and Web GUI server on the remote server. For more information about stopping IBM Spectrum Control servers use the following link.	1246
BPCIN0252E Db2 has been installed by using a domain user account. The IBM Spectrum Control installation software does not allow the Db2 user name userID because this user name exists in both the Windows domain and the local operating system repositories. Specify a Db2 user name that exists only in the Windows Domain registry.	1247
BPCIN0253W All IBM Spectrum Control components are installed. GUIErrorMessage	1247
BPCIN0254W IBM Spectrum Control servers are installed on remoteHost.	1247
BPCIN0255E The database dbName must be upgraded to the current version. For more information about upgrading Db2, go to the IBM Knowledge Center and search for "Upgrading Db2"	1247
BPCIN0256W If you continue upgrading to IBM Spectrum Control, the reports you have from Tivoli Storage Productivity Center Version oldVersion will be deleted. When you uninstall Tivoli Integrated Portal, the Authentication Services Server is also uninstalled. If you have storage subsystems that are configured to use LDAP authentication through the Authentication Services Server, before you upgrade the product and uninstall Tivoli Integrated Portal, reconfigure the storage subsystems so that these subsystems do not use LDAP authentication. Click OK to continue or Cancel to select another option.	1247
BPCIN0257E The Db2 database installation on dbPath path was made when the creation of 8.3 filenames was disabled on this server. The IBM Spectrum Control installation will fail when Db2 is installed to a path with spaces and 8.3 filenames were not enabled when it was installed.	1248
BPCIN0258E The LDAP configuration export did not succeed. For more information, review the log files, and in the LDAP export command output, search for "exportLDAPRepositories".	1248
BPCIN0260E The connection to the local database failed because the user name or password is not valid.	1248
BPCIN0261E The installation program was unable to connect to the local database. Check the server to verify that you can connect to Db2. For more information about Db2 connection issues, go to the IBM Knowledge Center at http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome .	1248
BPCIN0262E The IBM Spectrum Control upgrade process has stopped because Jazz for Service Management and Tivoli Integrated Portal are using the same ports. To continue, you must install Jazz for Service Management and Tivoli Common Reporting again and use ports that are different from the ports that are used by Tivoli Integrated Portal.	1249
BPCIN0264E The version of the database repository component cannot be queried from the Db2 Database Manager. SQL Error Message is sqlerrmsg	1249
BPCIN0265E An error occurred during the prevalidation of the component because the file_name is missing.	1249
BPCIN0266E Db2 has been installed by using a domain user name. The user name dbuser is not a member of the domain Db2 administrator group domainDB2AdminGroup, so the IBM Spectrum Control installation program cannot start.	1249
BPCIN0267E The password password cannot start with the following special characters: characters.	1250
BPCIN0268E The installation program could not stop the .You must manually stop the by running the script.	1250
BPCIN0269E The Jazz(tm) for Service Management installation image Version "retrivedVersion" is not at the correct build level.	1250
BPCIN0272E "installDirectory" cannot be found where the Jazz(tm) for Service Management installation files are extracted. Extract the Jazz(tm) for Service Management installation files to a local directory before you start the Jazz(tm) for Service Management installation program.	1250
BPCIN0275E A 32-bit version of Db2 has been detected. IBM Spectrum Control can only be installed with a 64-bit version of Db2. Install a supported 64-bit version of Db2, and install IBM Spectrum Control again.	1251
BPCIN0276E A required system library could not be loaded. Review the log files for more information about this error.	1251
BPCIN0277W The Jazz(tm) for Service Management installation image that you selected is not the latest version. Click Yes to continue with the current version and not install or upgrade reports. Click No to install or upgrade to Jazz(tm) for Service Management Version "latest_JazzSM_available_version" and Tivoli Common Reporting Version "latest_TCR_available_version".	1251
BPCIN0279W IBM Spectrum Control does not support the version of Jazz(tm) for Service Management that is installed on this computer. You must upgrade to Jazz(tm) for Service Management minJazzSMVersion before you can continue.	1251
BPCIN0281E The current version of the Tivoli Common Reporting installation image is not at the correct level.	1252
BPCIN0282E The installation package for Tivoli Common Reporting Version 3.1.0.1 cannot be found on your computer. To continue installing IBM Spectrum Control, download the installation package for Tivoli Common Reporting Version 3.1.0.1 in the same directory where you downloaded the installation package for Tivoli Common Reporting Version 3.1.0.2.	1252
BPCIN0283E You cannot upgrade "latest_version" to the lower version "lower_version".	1252
BPCIN0284E The Jazz for Service Management installation directory does not have execution rights. You can add execution rights to the directory by running the following command:chmod -R u+x "JazzSM_build_folder"	1252
BPCIN0285E The getDB2Inst.sh "user_name" command is displaying a "null" result because the Db2 profile was not sourced for user "user_name" or the environment for "user_name" has been corrupted.	1253
BPCIN0286W The operation of the storage resource agent is limited on the operating system of the local server. Total Disk Space and Available Disk Space on this server cannot be determined.	1253
BPCIN0287E Directory directory is located on a memory based file system (RAM disk) and cannot be used for installing IBM Spectrum Control.	1253
BPCIN0288E The upgrade process detected an error with the previously installed version of the product. The installation directory of the previously installed version of the product contains corrupted files. Remove or fix the corrupted files and run the upgrade process again. See the lax*-out.txt and lax*-err.txt log files in the system temporary directory for details of the error.	1253
BPCIN0289E You must install Db2 before you can install Cognos BI Reports. Ensure that Db2 is already installed on the system and the db2profile is sourced.	1254
BPCIN0290E The IBM Spectrum Control installation image is corrupted. Extract the IBM Spectrum Control installation image again into an empty directory.	1254
BPCIN0291E You need to define 'localhost' in the 'hosts' file that is used by Jazz for Service Management and restart the host system.	1254
BPCIN0293E The IBM Spectrum Control installation program cannot find the license folder. Ensure the license folder is extracted into "extractorDir".	1254
BPCIN0295E Tivoli Storage Productivity Center old_version is installed. Before you upgrade to IBM Spectrum Control version new_version, you must upgrade to Tivoli Storage Productivity Center version 5.2.7.	1255
BPCIN0296E The installation program could not rename the short_name directory because the folder is in use by other processes. Quit the installation program and stop all processes that access the directory full_path.	1255
BPCIN0297E The user_name user does not have the full control permission for the folder_path folder.	1255
BPCIN0298E The user_name fenced user does not have full permissions for the directory_path database directory.	1255
BPCIN0299E The installation location "install_location" has a naming conflict with a file or folder named folder_path. Rename the folder_path file or folder.	1255
BPCIN0300E The value db2_variable_existing_value is not allowed for the Db2 variable DB2_LIMIT_FENCED_GROUP. Use the db2set -g DB2_LIMIT_FENCED_GROUP=OFF command to change the value of the variable to OFF.	1256

BPCIN0301E The versionFile file is empty. Extract IBM Spectrum Control again and make sure the version.txt file contains a valid build string.	1256
BPCIN0302E Db2 Advanced Enterprise Server Edition is not supported. Install Db2 Enterprise Server Edition and then proceed with your installation.	1256
BPCIN0303W Your storage environment has switches with obsolete data sources. Configure up-to-date data sources after you upgrade.	1256
BPCIN0304E You cannot upgrade from a Basic Edition license.	1257
BPCIN0306W The Storage Resource agent registered successfully with the Data server but some problems occurred after the registration. Review the log files in the SRA_log_name directory for additional information.	1257
BPCIN0307W You are currently connected to devices using a CIM interface. This upgrade might require new certificates to be generated on your CIM managed devices to continue monitoring them. To resolve this issue, go to http://www.ibm.com/support/docview.wss?uid=swg21976237	1257
BPCIN0307E The security certificates for the Web server have expired. Renew the security certificates and run the installation program again. See the lax*-out.txt and lax*-err.txt log files in the system temporary directory for details of the error and how to resolve the issue.	1257
BPCIN0308W Brocade switches that use SNMP services were detected. SNMP data sources are no longer used to manage Brocade fabrics and switches. Instead, you must use the embedded SMI agent in Brocade Network Advisor.	1258
BPCIN0309E The port configuration export did not succeed. For more information, review the log files, and in the port export command output, search for "exportCurrentPorts".	1258
BPCIN0310E Spectrum Control installation requires that the Administrators group in Windows is assigned "Debug programs" privilege. Check this setting, log out of and back into Windows, and try again.	1258
BPCIN0311E Spectrum Control installation requires the wmic command to work. Check this command, ensure the service "Windows Management Instrumentation" is running and try again.	1259
BPCIN0312E Spectrum Control installation requires that wmic command to work. Check this command, ensure service "Windows Management Instrumentation" is working, disable any antivirus and try again.	1259
BPCIN0313E Spectrum Control installation requires that wmic command to work. Check this command and the PATH environment variable and try again.	1259
BPCIN0314E Spectrum Control installation requires that wmic command to work. Check this command and try again.	1259
BPCIN0315E Spectrum Control installation requires that chcp command to work. Check this command and try again.	1260
BPCIN0316E The IBM Spectrum Control upgrade process does not support custom materialized query tables (MQTs). Remove any MQTs from the database and run the upgrade process again.	1260
BPCIN0317E The IBM Spectrum Control upgrade process does not support missing vendor information from the .com.zerog.registry.xml file. Add the missing information and run the upgrade process again. See the lax*-out.txt and lax*-err.txt log files to resolve the issue.	1260
BPCIN0318E The IBM Spectrum Control upgrade process does not support duplicated users user in IBM WebSphere Application Server related to the Web server. Remove the duplicate user from a federated repository other than the localOS default repository; run the upgrade process again.	1260
BPCIN0319E The currently installed version of AIX XL C/C++ RUNTIME is not supported. You must upgrade to a supported level. The version that was detected is version.	1260
BPCIN0320E PAM (Pluggable Authentication Modules) isn't installed on your system. Installing PAM would resolve the issue.	1261
BPCIN0321E An error occurred when restoring authorities before creating the database dbName. Review the log files for more information.	1261
BPCIN0323E The Data Server port has different values in the configuration files. Verify that the port value is correct in the installDir/data/config/server.config file, and the installDir/config/InstallVariable.properties file. Also, verify that the PORT_NUMBER column where SERVER_TYPE is serverType is correct in the T_RES_Server database table.	1261
BPCIN0324W The current certificate used to secure the connection of the IBM Spectrum Control Data server and the Storage Resource agents cannot be replaced with the latest certificate that provides higher security. After you complete the upgrade, you can create new certificates by following the information in the following link.	1262
BPCIN0325E IBM Tivoli Storage Productivity Center old_version is installed. Before you can upgrade to IBM Spectrum Control new_version, you must upgrade to IBM Spectrum Control 5.3.0 or later. For more information about upgrading to IBM Spectrum Control 5.4.0 or later, go to the following link:	1262
BPCIN0328E The IBM Spectrum Control installation requires the LD_LIBRARY_PATH environment variable be set. Set this environment variable and try the installation again. If you ran setup.bin from the IBM Spectrum Control installation directory with sudo and the LD_LIBRARY_PATH variable was properly set, then try the sudo -E LD_LIBRARY_PATH=\$LD_LIBRARY_PATH ./setup.bin command.	1262
BPCIN0329E The IBM Spectrum Control installation requires that the LIBPATH environment variable be set. Set this environment variable and try the installation again.	1263
BPCIN0330E An SQL exception was created when querying the Db2 Database Manager. The SQL error message is: sqlerrmsg.	1263
BPCIN0331E The tpcregFile file is not valid because it contains a failed or partial upgrade to a previous IBM Spectrum Control Version newTPCVersion. Continue with the incomplete upgrade before you try to upgrade to the new IBM Spectrum Control Version varNewUpgradeVersion.	1263
BPCIN0335E The IBM Spectrum Control installation requires you install the libstdc++ package on the AIX operating system. Download and install the libstdc++ package and start the installation again. You can download the package here: https://www.ibm.com/developerworks/aix/library/aix-toolbox/alpha.html .	1263
BPCIN0336E The version of the libstdc++ package you installed is lower than the minimum required Version minVersion. Download and upgrade the libstdc++ package and start the installation again. You can download the package here: https://www.ibm.com/developerworks/aix/library/aix-toolbox/alpha.html .	1264
BPCIN0337E The componentServer password failed to validate. Check that you entered the password correctly and try again.	1264
BPCIN0338E The componentServers passwords failed to validate. Check that you entered the passwords correctly and try again.	1264
BPCIN0339W Your current Db2 instance is using a trial license. The license will expire on expireDate. Once your trial expires Db2 will not start. If you want to upgrade your Db2 license now, use the information in the following link.	1264
BPCIN0340W Your current Db2 instance is using a trial license. The license will expire on expireDate. Once your trial expires Db2 will not start. If you want to upgrade your Db2 license now, goto link.	1264
BPCIN0343W There are new certificate requirements that are strictly enforced for macOS Catalina users which might affect their ability to access the IBM Spectrum Control GUI. During an upgrade of IBM Spectrum Control, certificates self-signed by IBM Spectrum Control will be made compliant automatically. However, if one or more of your certificates are not self-signed by IBM Spectrum Control, see the following links for more information. Validate that your certificates are compliant.	1265
BPCIN0344E Your current Db2 instance trial license has expired. If you want to upgrade your Db2 license now, goto the following link.	1265
BPCIN0345E The user user is part of the "Deny access to this computer from the network" security policy setting on your computer. Either contact your administrator to have the user removed from the security policy or enter a different user. For more information on required user privileges in installation scenarios, go to the following link.	1265
BPCIN0346E Your current Db2 License type: Community is not supported for IBM Spectrum Control. To upgrade your Db2 license, go to the following link:	1266
BPCRE - Alert server messages	1266
BPCRE0001I Trying again to connect to the repository database...	1266
BPCRE0003E An invalid type was specified for a day range.	1266
BPCRE0007E Error adding rule rule_name.	1266
BPCRE0008I Drools Version: version_number.	1267
BPCRE0009I Duration - build rules:value seconds.	1267
BPCRE0010I Duration - setup KIE session: value seconds.	1267

BPCRE0011E Failure writing rule rule to database.	1267
BPCRE0012E Failure inserting data for rule rule to database.	1267
BPCRE0013E Failure rebuilding rules after membership change in group.	1267
BPCRS - Spectrum Control Middleware messages	1268
BPCRS0000E A required key key is not found in map.	1268
BPCRS0001E Value value is not found in list.	1268
BPCRS0002E Some or all of the required key-value bindings are not found. Required Keys: required. Supplied keys: supplied.	1268
BPCRS0003E key value value is not in list.	1268
BPCRS0004E Expected non-null value for var.	1268
BPCRS0005E Expected non-empty string value for str.	1269
BPCRS0006E Value supplied for name is not of type type.	1269
BPCRS0007I The log file retention settings were successfully updated.	1269
BPCRS0008E The log file retention settings were not successfully updated.	1269
BPCRS0009I The alert disposition settings were updated.	1269
BPCRS0010E The alert disposition settings were not updated.	1269
BPCRS0011I The history and log retention settings were successfully updated.	1270
BPCRS0012E The history and log retention settings were not successfully updated.	1270
BPCRS0013I Step currentStep of totalSteps : stepName	1270
BPCRS0014I stepName completed	1270
BPCRS0015E stepName failed	1270
BPCCM - Data collector messages	1270
BPCCM0001I Data collection is being performed by hostname.	1270
BPCDP - Data processor messages	1271
BPCDP0000I Performance data for natural key resource at date and time timestamp was collected and processed successfully.	1271
BPCDP0001E Error while collecting and processing performance data for natural key resource at date and time timestamp. Performance data was not collected and processed.	1271
BPCDP0002E The processing of performance data for the resource could not be completed.	1271
BPCDP0003E No performance data is available at the current time for this resource.	1272
BPCDP0004I Performance data was retrieved and persisted but aggregation of data to higher-level components didn't complete because the relationship to the higher-level components couldn't be determined.	1272
BPCDP0005E Could not save the performance data that was collected from the resource.	1272
BPCDP0006I Performance data at date and time timestamp was processed and saved successfully for the resource.	1272
BPCDP0007E The resource is missing. The resource is required.	1272
BPCDP0008E Identifying information for the resource is missing. This information is required.	1272
BPCDP0009E Information identifying the resource type is invalid: system type .	1273
BPCDP0010E Information uniquely identifying the resource is missing. This information is required.	1273
BPCDP0011E Information uniquely identifying the resource is invalid.	1273
BPCDP0012E The UUID for the tenant's resource is invalid.	1273
BPCDP0013E The start time for the performance data is invalid: start time	1273
BPCDP0014E The end time for the performance data is invalid: end time	1274
BPCDP0015I Performance data at date and time timestamp was processed and saved successfully for the resource, but the data processing raised warnings.	1274
BPCDP0016W Performance data for natural key resource at date and time timestamp was collected and processed successfully, but the data processing raised warnings.	1274
BPCCS - Scheduler messages	1274
BPCCS0000E An error occurred while collecting performance data from the device. The collection is being attempted by a different collector.	1275
BPCCS0001W The data collection is taking longer than expected.	1276
BPCCS0002E Currently, there is no data collector available for this device.	1276
BPCCS0005I Performance monitor is starting at an interval of interval interval units. This action was requested by user name.	1276
BPCCS0008I Collection can no longer continue due to invalid credentials. Use the 'Modify Connection' dialog to fix the storage system credentials and resume collection.	1276
BPCCS0009E Failed to save the performance monitor schedule.	1276
BPCCS0010E A job cannot be run for resource resourceName because there is a job already running for the resource.	1277
BPCCS0011W The schedule change was saved but the update to the active collection did not happen.	1277
BPCCS0012I Performance monitor is stopped. This action was requested by user name.	1277
BPCCS0013I Performance monitor is stopped.	1277
BPCCS0014I Performance monitor is starting at an interval of interval interval units.	1278
BPCCS0015I Performance monitor collection interval was updated to interval interval units. This action was requested by user name.	1278
BPCCS0016I Performance monitor collection interval was updated to interval interval units.	1278
BPCCS0017I Performance monitor is enabled. This action was requested by user name.	1278
BPCCS0018I Performance monitor is enabled.	1278
BPCCS0019I Performance monitor is disabled. This action was requested by user name.	1278
BPCCS0020I Performance monitor is disabled.	1278
BPCCS0021W Performance monitor is starting. The initial attempt to start collection failed so it is retried. This action was requested by user name.	1278
BPCCS0022W Performance monitor is starting. The initial attempt to start collection failed so it is retried.	1279
BPCCS0023I Performance monitor collection interval is enabled and updated to interval interval units. This action was requested by user name.	1279
BPCCS0024I Performance monitor collection interval is enabled and updated to interval interval units.	1279
BPCCS0025E Access to the agent or device is denied. Ensure that valid credentials are specified for agent agent name.	1279
BPCCS0026E New performance data is not yet available for the device. Statistics with time stamps later than time stamp could not be found.	1279
BPCCS0027E The performance monitor failed due to an internal error.	1280
BPCCS0028E The value that is specified as parameter (value) is invalid.	1280
BPCCS0029E Cannot connect to the device with the address IP address.	1280
BPCCS0030E Cannot connect to the SNMP data source IP address.	1281

BPCSS0031E Cannot authenticate with the provided user credentials.	1281
BPCSS0032E Passphrase is incorrect for subsystem param1.	1281
BPCSS0033E Passphrase is required. Specify one for subsystem param1.	1281
BPCSS0034E Verify that they private key that was provided for subsystem param1 was in the OpenSSH file format. If it is in another format, it needs to be converted before it can be used.	1281
BPCSS0035E The user does not have the required authority to complete the task or command.	1282
BPCSS0036E Cannot connect to the storage system or cluster.	1282
BPCSS0037W The device cannot be reached.	1282
BPCSS0038E The device or device agent did not respond within the allotted time.	1282
BPCSS0039E The host name or IP address {0} is not valid.	1283
BPCSS0040E The host name or IP address is not valid.	1283
BPCSS0041E Cannot connect to the device.	1283
BPCSS0042E Cannot connect to the SNMP data source.	1283
BPCSS0043E Passphrase is incorrect.	1284
BPCSS0044E Passphrase is required.	1284
BPCSS0045E Access to the device is denied. Ensure that valid credentials are specified.	1284
BPCSS0046E Verify that they private key that was provided was in the OpenSSH file format. If it is in another format, it needs to be converted before it can be used.	1284
BPCSS0047E New performance data is not yet available for the device.	1284
BPCSS0048E The parameter for the Performance Manager API is invalid.	1285
BPCSS0049E Schedule is not enabled for the resource resource.	1285
BPCSS0050W Performance data could not be collected for device device name because the device or data source cannot be reached (reason reason code). The current samples are skipped.	1285
BPCSS0051E The device or device agent did not respond within the allotted time (timeout valuesseconds).	1286
BPCSS0052W Performance data continuity is broken. The device might have been reset or rebooted. record count performance data records were discarded.	1286
BPCSS0053W No valid performance data was provided by the monitored resource. Zero performance data records were inserted into the database.	1287
BPCSS0054E A timeout occurred while polling the performance statistics for this device: device name	1287
BPCSS0055E Performance data was not collected for device device name due to error error trace. The current samples are skipped.	1287
BPCSS0056E The last performance Data Collection was not readable for device device name, the collection failed with error error trace.	1287
BPCSS0057E Cannot connect to the switch with the provided IP address, host name, protocol, and port.	1287
BPCSS0058E Cannot authenticate to the switch with the provided user name and password.	1288
BPCSS0059E The specified user name does not have the required permissions for the switch.	1288
BPCSS0112I The probe was stopped. This action was requested by user name.	1288
BPCSS0113I The probe is stopped.	1288
BPCSS0114I The probe is starting at an interval of interval interval units.	1289
BPCSS0115I The probe interval was updated to interval. This change was requested by user name.	1289
BPCSS0116I The probe interval was updated to interval.	1289
BPCSS0117I The probe is enabled. This action was requested by user name.	1289
BPCSS0118I The probe is enabled.	1289
BPCSS0119I The probe is disabled. This action was requested by user name.	1289
BPCSS0120I The probe is disabled.	1289
BPCSS0123I The Probe collection interval is enabled and updated to interval. This action was requested by user name.	1289
BPCSS0124I Probe interval is enabled and updated to interval.	1290
BPCSS0105I Probe is starting at an interval of interval interval units. This action was requested by user name.	1290
BPCUI - User Interface messages	1290
BPCUI0000E The action can't be completed because the following error occurred: Error message text.	1297
BPCUI0001E An action could not be completed and the following error message was generated: TPCRemoteException message	1297
BPCUI0002E Failed to retrieve the requested data because the service is unavailable.	1298
BPCUI0003E The NAPI with the IP address Napi IP was not added because of an Internal Error	1298
BPCUI0004E The SSH private key for the NAPI Napi IP could not be uploaded	1298
BPCUI0005E The action cannot be completed because the following internal error has occurred: message.	1298
BPCUI0007E The discovery job failed to complete.	1298
BPCUI0009E The SSH key could not be loaded for the following reason:IOException message	1299
BPCUI0010E The host name or IP address that you entered is a resource_type, but you selected to add a different type of storage system.	1299
BPCUI0011E The Device Server did not discover any device	1299
BPCUI0012E Cannot connect to the device with the address Ip Address.	1299
BPCUI0019E No data is available for this selection.	1299
BPCUI0025E Probe job job Id failed.	1299
BPCUI0029E Invalid parameter param passed.	1300
BPCUI0030I This task was already executed.	1300
BPCUI0032E An unexpected response was received from the server.	1300
BPCUI0034E Invalid number of runs to keep for each schedule. The number should be between param1 and param2.	1300
BPCUI0035E Invalid number of days' worth of log-files to keep. The number should be between param1 and param2.	1300
BPCUI0036E The schedule id scheduleID associated with this job is no longer valid. It might have been deleted. Refresh the view and try again.	1301
BPCUI0037E The replication server is not installed or is unavailable.	1301
BPCUI0038E Invalid number of days to retain alerts. The number should be between param1 and param2.	1301
BPCUI0039E A Storage Resource agent cannot be found.	1301
BPCUI0040E Parsing results from a call to the Data server failed with the following error message: param1.	1301
BPCUI0042E Communication with the Data Server failed with the following error: param1	1302
BPCUI0043E Cannot connect to the Data server.	1302
BPCUI0044E The entity was not found in the database.	1302
BPCUI0045E Host name length exceeds the 255 character limit	1302

BPCUI0046E Report 'configurationId' not found	1302
BPCUI0047E Parameter 'parameterName' is not defined in report configurationId'	1302
BPCUI0048E No property is not defined for report configurationId'	1303
BPCUI0049E No such property propertyName for report configurationId'	1303
BPCUI0050E variableName can not be overridden	1303
BPCUI0051E variableName not valid report output format.	1303
BPCUI0052E variableName not reachable	1303
BPCUI0053E Cannot authenticate with the provided user credentials.	1303
BPCUI0054E The host name or IP address {0} is not valid.	1304
BPCUI0055E Cannot connect to the storage system.	1304
BPCUI0056E Cannot connect to the storage system or cluster.	1304
BPCUI0058I No supported resources were discovered on the data source data_Source_Address.	1304
BPCUI0060I File param was successfully uploaded to the Data Server.	1304
BPCUI0061E Upload file type param is not supported.	1305
BPCUI0062E The requested action failed with the following error message: error message	1305
BPCUI0063E Cannot find jobs for scheduleId param and deviceId param. No logs are displayed.	1305
BPCUI0064E A log file cannot be displayed for the job.	1305
BPCUI0065E The job log file cannot be accessed. The log file may have been manually removed or may have been deleted because it was older than retain_days days or it exceeded the maximum number of no_of_lofs runs.	1305
BPCUI0067E The schedule for collecting status and asset data cannot be created.	1306
BPCUI0068E A proposed schedule for collecting status and asset data cannot be created.	1306
BPCUI0069E The proposed schedule for collecting status and asset data cannot be deleted.	1306
BPCUI0071E The task task_name could not be completed.	1306
BPCUI0072E Cannot connect to the Device server. Verify that the database service and Device server are running, and that the Device server is accessible.	1307
BPCUI0073E Can't make a connection to the storage_resource storage resource.	1307
BPCUI0074E The wizard could not set an attribute for the storage resource.	1307
BPCUI0075E The certificate wasn't saved on the server.	1307
BPCUI0076W The initial job to collect status and asset data did not start.	1307
BPCUI0077E A failure occurred loading the certificate.	1308
BPCUI0078I The certificate was loaded successfully.	1308
BPCUI0079E The SSL certificate is not in the expected format.	1308
BPCUI0084W The wizard could not retrieve the default interval information for performance monitoring.	1308
BPCUI0085E The user name or password for the hypervisor or vCenter hypervisor or vCenter Server is invalid.	1308
BPCUI0086E The SSL certificate is invalid for the hypervisor or vCenter hypervisor or vCenter Server, or the firewall is blocking access to it.	1309
BPCUI0087E The version of the hypervisor or vCenter hypervisor or vCenter Server is not supported.	1309
BPCUI0088E The host name, protocol, or port for the hypervisor or vCenter hypervisor or vCenter Server is invalid, or the hypervisor or vCenter Server is unreachable.	1309
BPCUI0089W Cannot retrieve a valid set of data collection intervals for performance monitoring.	1310
BPCUI0090I All alerts were removed.	1310
BPCUI0091W error_count of total_count alerts were not removed.	1310
BPCUI0093I No data path is available for deviceNameVariable.	1310
BPCUI0094E Authorization failed due to an internal error.	1310
BPCUI0097E Authorization failed due to an invalid request context.	1311
BPCUI0098E The current user is not authorized to perform the requested function.	1311
BPCUI0099E The storage resource is not available.	1311
BPCUI0100I success_count alerts were marked as acknowledged.	1311
BPCUI0101I The alert was marked as acknowledged.	1312
BPCUI0102E None of the alerts were marked as acknowledged.	1312
BPCUI0104I success_count alerts were marked as unacknowledged.	1312
BPCUI0105I The alert was marked as unacknowledged.	1312
BPCUI0108I All informational alerts were marked as acknowledged.	1312
BPCUI0110W Some informational alerts were not marked as acknowledged.	1312
BPCUI0111I All alerts were marked as acknowledged.	1313
BPCUI0112I success_count alerts were removed.	1313
BPCUI0113I The alert was removed.	1313
BPCUI0114I All acknowledged alerts were removed.	1313
BPCUI0116W Some acknowledged alerts were not removed.	1313
BPCUI0120W Some acknowledged alerts were not marked as unacknowledged.	1314
BPCUI0121E Unable to communicate with the product server. Make sure that the server is running properly.	1314
BPCUI0122E No job log file was created for this job run.	1314
BPCUI0123E The action cannot be completed.	1314
BPCUI0124E An unexpected error occurred during the execution of the action.	1315
BPCUI0125E The alert is not available.	1315
BPCUI0126E The status of the Performance Monitors could not be retrieved.	1315
BPCUI0127E The currently installed version of the product does not have the required product license for the function that you requested.	1315
BPCUI0128E An undefined capacity chart metric was requested.	1315
BPCUI0129I Alerts that were migrated from a previous version of the product are not shown on this page.	1316
BPCUI0130E The alerts cannot be acknowledged because they were deleted.	1316
BPCUI0131E The alerts cannot be unacknowledged because they were deleted.	1316
BPCUI0132W success_count alerts were marked as acknowledged. unsuccess_count alerts cannot be marked as acknowledged because they were deleted.	1316
BPCUI0133W success_count alerts were marked as unacknowledged. unsuccess_count alerts cannot be marked as unacknowledged because they were deleted.	1316
BPCUI0134E The alert cannot be acknowledged because it was deleted.	1317

BPCUI0135E The alert cannot be unacknowledged because it was deleted.	1317
BPCUI0136E The device was not removed because the action is not supported for devices of type devType.	1317
BPCUI0137E Input text provided has invalid character(s): characters. Input text: text	1317
BPCUI0141E Host name or IP address hostname specified on line line of file file is not valid.	1318
BPCUI0143E Host port WWPN wwpn specified on line line of file file is not valid.	1318
BPCUI0144E Duplicate server name specified on lines line1 and line2 of file file.	1318
BPCUI0145E Could not parse file file.	1318
BPCUI0146E Could not parse file file. Invalid entry on line line.	1318
BPCUI0148I Successfully deleted server server_name.	1318
BPCUI0149I Successfully modified ports of server server_name.	1319
BPCUI0150I The server was created.	1319
BPCUI0151E The host name or IP address is associated with another resource.	1319
BPCUI0152I The data source data_Source_Address was successfully added as a data source for monitoring. The following new resources were detected:	1319
BPCUI0155W You cannot provision volumes because there is no Fibre Channel host port information for at least one server.	1319
BPCUI0156W You cannot provision volumes to servers that use different operating systems.	1320
BPCUI0157W You cannot provision volumes to servers and virtual machines at the same time. To provision volumes, ensure that you select either only servers or only virtual machines.	1320
BPCUI0158I Volumes are assigned to the hypervisors that host virtual machines. Volumes are not assigned directly to virtual machines.	1320
BPCUI0159W You cannot provision volumes because at least one of the hypervisors that host the virtual machines is not being monitored. Ensure that all the hypervisors that are hosting the virtual machines that were selected for provisioning were probed.	1320
BPCUI0160E Duplicate port WWPN wwpn specified on lines line1 and line2 of file file.	1321
BPCUI0162W File file does not contain any servers to create.	1321
BPCUI0166W Optimization cannot be done in place to the subsystem since storage subsystem param1 and/or its pools belong to more than one capacity pool. Following are capacity pools the subsystem is associated with: param2	1321
BPCUI0167W Optimization cannot be done in place to the subsystem since storage subsystem param1 and/or its pools are not part of any capacity pool.	1321
BPCUI0168W Optimization cannot be done in place to the server param1 since storage subsystems or storage pools associated with luns assigned to the server belong to more than one capacity pool. Following are associated capacity pools: param2	1321
BPCUI0169W Optimization cannot be done in place to the server param1 since storage subsystems or storage pools associated with luns assigned to the server are not part of any capacity pool.	1322
BPCUI0170W Optimization cannot be done in place to the storage entity param1 since storage subsystems or storage pools associated with it belong to more than one capacity pool. Following are associated capacity pools: param2	1322
BPCUI0171W Optimization cannot be done in place to the storage entity param1 since storage subsystems or storage pools associated with it are not part of any capacity pool.	1322
BPCUI0172E The operation timed out while waiting for a response from the server.	1322
BPCUI0173E File file does not exist or is empty.	1323
BPCUI0174E The device does not support the credential mechanism used.	1323
BPCUI0175E A required parameter is missing.	1323
BPCUI0176E The highlighted field contains an invalid value.	1323
BPCUI0177E The highlighted field contains a value that is outside of the allowed range. The value must be between minVal and maxVal.	1323
BPCUI0178E A service class with the same name and type already exists.	1323
BPCUI0179I The service class was created.	1324
BPCUI0180I Based on the known configuration of storage system host connections, fabric zone aliases, and HBA ports, additional ports may have been added to the selection below.	1324
BPCUI0181I You selected to add a expectedDevice resource, but a foundDevice resource was detected and will be added.	1324
BPCUI0182I The data source data_Source_Address was added as a data source for monitoring. No new resources were detected.	1324
BPCUI0183E The text in the highlighted field exceeds the maxLength character limit.	1325
BPCUI0185W Unable to lookup the IP Address for Host Name hostName. Enter the IP Address manually.	1325
BPCUI0189I Configuration of SRA deployment and probe schedules were done successfully.	1325
BPCUI0190W Configuration of SRA finished with some warnings or errors. Check the detail messages.	1325
BPCUI0191E An internal error occurred while testing connecton to param1.	1325
BPCUI0192E The supplied service class type is invalid.	1326
BPCUI0193E The specified SMI-S provider was not found. Make sure that the protocol, SMI-S provider host name or IP address, and port are specified correctly and that the SMI-S provider is properly configured at that location.	1326
BPCUI0194E An unknown error has occurred. Please review all values entered.	1326
BPCUI0195E The Interop Namespace is not correct. Please correct this entry.	1326
BPCUI0196E A timeout ocured while processing the request. Please retry request.	1326
BPCUI0197E A connection was not established. Make sure that the protocol, SMI-S provider host name or IP address, and port are specified correctly.	1326
BPCUI0198E The authentication to the SMI-S provider failed.	1327
BPCUI0199E An SSLHandshakeException or SSLProtocolException has occurred. This exception might be due to an invalid SLP registration, e.g. 'http' instead of 'https'.	1327
BPCUI0201E There is a pending delete in process for this SMI-S provider.	1327
BPCUI0202I Success	1327
BPCUI0203E The selected resources were not removed.	1327
BPCUI0204W successfulDeletes of attemptedDeletes of the selected resources were removed.	1327
BPCUI0205W successfulDeletes selected resources were removed, however warnings did occur.	1327
BPCUI0209E A database operation cannot be completed.	1328
BPCUI0210I Device param1 supports performance monitoring.	1328
BPCUI0211E No performance data is available for a resource.	1328
BPCUI0212E There is no Secure Shell running at this host/IP.	1328
BPCUI0213E Unsupported Secure Shell protocol was used.	1328
BPCUI0214E Invalid public key location for subsystem param1.	1329
BPCUI0215E Invalid public key format for subsystem param1.	1329
BPCUI0216E Passphrase was incorrect for subsystem param1.	1329

BPCUI0217E Unable to transfer the key(s) to the server param1.	1329
BPCUI0218E The specified private key file format is not supported. Please convert it to Open SSH (.pem) key file format for subsystem param1.	1329
BPCUI0219E The specified key file or key file name is already linked to another user.	1329
BPCUI0220E The IP address that was entered was the address of the management console for the storage system. You must enter the valid IP address of the block component of the storage system.	1329
BPCUI0221E The IP address you entered is the address of another device's management console.	1330
BPCUI0222E The IP address you entered points to a device of another type.	1330
BPCUI0223E Passphrase is required. Specify one for subsystem param1.	1330
BPCUI0224E Cannot connect to a resource because of an SSL certificate error. Troubleshooting information: http://www.ibm.com/support/docview.wss?uid=swg21976237	1330
BPCUI0225I The agent log files for server_Name have been collected and copied to log_Location.	1331
BPCUI0226I Discovery of data_source is taking longer than expected. Click Close to run the discovery in the background.	1331
BPCUI0227E Thin provisioning must be enabled when compression is enabled.	1331
BPCUI0229I 1 resource was added to name.	1331
BPCUI0231I count resources were added to name.	1331
BPCUI0233E The specified host name is already associated with an existing server.	1331
BPCUI0234E The specified IP address is already associated with an existing server.	1332
BPCUI0235E The specified host name and IP address are already associated with an existing server.	1332
BPCUI0236E The disabling of the agents failed.	1332
BPCUI0237E Errors occurred when attempting to disable some of the agents.	1332
BPCUI0238W Warnings occurred when attempting to disable warningCount of the agents.	1332
BPCUI0239I attemptedCount of the selectedCount selected agents were disabled.	1333
BPCUI0240E The agents were not enabled.	1333
BPCUI0241E Errors occurred when attempting to enable some of the agents.	1333
BPCUI0242W Warnings occurred when attempting to enable warningCount of the agents.	1333
BPCUI0243I attemptedCount of the selectedCount selected agents were enabled.	1333
BPCUI0244I The credentials of an agent were updated.	1334
BPCUI0245I The credentials of updateCount agents were updated.	1334
BPCUI0246E Cannot authenticate to the file module with the provided user credentials.	1334
BPCUI0247E Unknown file module key user.	1334
BPCUI0248E The SSH key could not be loaded for the following reason:IOException message	1334
BPCUI0249E Passphrase is incorrect.	1334
BPCUI0250E Passphrase is required.	1335
BPCUI0251E Cannot connect to the storage system or cluster.	1335
BPCUI0252E The host name or IP address {0} is not valid.	1335
BPCUI0253E Cannot connect to the data source for the resource with the address ip_address.	1335
BPCUI0254E Invalid private key location.	1335
BPCUI0255W The following resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?	1335
BPCUI0256W The following resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?	1336
BPCUI0257W The following resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?	1336
BPCUI0258W The following internal resources of a storage system you are attempting to add are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?	1336
BPCUI0259W The following storage systems and storage-system internal resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?	1336
BPCUI0260E The specified private key file format for the file module is not supported. Please convert it to Open SSH (.pem) key file.	1336
BPCUI0261E The service class was not found in the database.	1337
BPCUI0262E The capacity pool was not found in the database.	1337
BPCUI0263E The scheduling of the agent upgrade jobs failed.	1337
BPCUI0264E Errors occurred when attempting to schedule the upgrade jobs of some of the agents.	1337
BPCUI0265W Warnings occurred when scheduling the upgrade of warningCount of the agents.	1337
BPCUI0266I attemptedCount of the selected agents were scheduled for upgrade.	1337
BPCUI0267I The upgrade agent job was successfully scheduled for hostName.	1338
BPCUI0268W Deleting a capacity pool does not affect any volumes or shares that were provisioned from the capacity pool. However, the volumes or shares are no longer associated with the capacity pool. Associations with the following volumes or shares will be removed:	1338
BPCUI0269W The following volumes are associated with the service class scName. When the volumes were created, they satisfied the requirements of the service class. If you modify the service class, the volumes are still associated with the service class, but might not satisfy the new requirements of the service class. Depending on your changes to the service class, users might incorrectly assume that the volumes have properties that they do not possess.	1338
BPCUI0270W The following shares are associated with the service class scName. When the shares were created, they satisfied the requirements of the service class. If you modify the service class, the shares are still associated with the service class, but might not satisfy the new requirements of the service class. Depending on your changes to the service class, users might incorrectly assume that the shares have properties that they do not possess.	1338
BPCUI0271W The following volumes are associated with the service class scName. If you delete the service class, the volumes are no longer associated with any service class.	1339
BPCUI0272W The following shares are associated with the service class scName. If you delete the service class, the shares are no longer associated with any service class.	1339
BPCUI0273E The action does not support the specified type of device.	1339
BPCUI0274I The connection test to resource data_Source_Name was successful.	1339
BPCUI0275I To collect data about zoning or complete zoning actions during provisioning, you must deploy Storage Resource agents to one or more servers that are on the fabric.	1339
BPCUI0276I Agent agentName was disabled.	1339
BPCUI0277I Agent agentName was enabled.	1340
BPCUI0278I The credentials for agentName were updated.	1340
BPCUI0279I There is no job defined for the device Name. Please create a job first before running it again.	1340
BPCUI0280I No switches are managed by the data_Source_Address data source.	1340

BPCUI0282I The resources that are managed by data_Source_Address are already known. One or more resources were added.	1340
BPCUI0284I No fabrics are managed by the data_Source_Address data source.	1340
BPCUI0286I The fabrics that are managed by data_Source_Address are already being monitored.	1341
BPCUI0289W The following network shared disks (NSDs) are already assigned to a capacity pool. Are you sure you want to move these NSDs to a different capacity pool?	1341
BPCUI0290W The following file systems and network shared disks (NSDs) are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?	1341
BPCUI0291W The following network shared disks (NSDs) are already assigned to a capacity pool. Are you sure you want to move these NSDs to a different capacity pool?	1341
BPCUI0292E The host name or IP address ip_address_or_hostname cannot be reached.	1341
BPCUI0293I A probe is started for deviceName.	1342
BPCUI0294I A performance monitor is started for deviceName.	1342
BPCUI0295I The performance monitor is stopped for deviceName.	1342
BPCUI0297W One resource was added to capacity_pool_name. One resource could not be added because it could not be found.	1342
BPCUI0298W count resources were added to capacity_pool_name. One resource could not be added because it could not be found.	1342
BPCUI0299W One resource was added to capacity_pool_name. count_Not_Found resources could not be added because they could not be found.	1342
BPCUI0300W count resources were added to capacity_pool_name. count_Not_Found resources could not be added because they could not be found.	1343
BPCUI0301E Failed to assign the role name role.	1343
BPCUI0302E Failed to retrieve the existing role assignments.	1343
BPCUI0303E Failed to remove all role assignments from the specified groups.	1343
BPCUI0304W An error occurred when saving the user-defined properties of the resourcesType.	1343
BPCUI0305E A capacity pool with the same name already exists.	1344
BPCUI0306W The selected resource was removed, however warnings did occur.	1344
BPCUI0307E The schedule could not be deleted.	1344
BPCUI0308I The resource does not have a connection configured. To add a connection to the resource, click Add Storage System.	1344
BPCUI0309I A probe schedule is defined for deviceName.	1344
BPCUI0310I A performance monitor schedule is defined for deviceName.	1344
BPCUI0311I Probe and performance monitor schedules are defined for deviceName.	1345
BPCUI0312I SNMP Discovery of switches is taking longer than expected. Click Close to run the discovery in the background.	1345
BPCUI0313I An upgrade is started for server deviceName.	1345
BPCUI0314E Failed to retrieve the list of user groups from the WebSphere user repository.	1345
BPCUI0315E Failed to retrieve the list of user groups from user repository due to an invalid search string.	1345
BPCUI0316W Failed to update the role cache maintained by the Device server.	1345
BPCUI0317E Access can not be removed, because at least one Administrator user must remain in the system.	1346
BPCUI0318E The group mapping can not be modified, because at least one Administrator user must remain in the system.	1346
BPCUI0319I A task is started for resource resourceName.	1346
BPCUI0320I Probe and performance monitor schedules are defined for deviceName. A performance monitor is scheduled to collect performance data after the probe is done.	1346
BPCUI0321I A task is paused for resource resourceName.	1346
BPCUI0322E A task could not be paused for resource resourceName.	1346
BPCUI0323I A task is resumed for resource resourceName.	1346
BPCUI0324E A task could not be resumed for resource resourceName.	1347
BPCUI0325E Failed to retrieve the list of users from the WebSphere user repository.	1347
BPCUI0326E Failed to retrieve the list of users from user repository due to an invalid search string.	1347
BPCUI0327E Failed to get the roles associated with the current user.	1347
BPCUI0328I A task is saved.	1347
BPCUI0329I A task was successfully deleted.	1347
BPCUI0330E The user user is not authorized to access the product.	1347
BPCUI0331I A task is cancelled for resource resourceName.	1348
BPCUI0332E An unexpected error occurred. The task for schedule schedule name could not be paused or resumed.	1348
BPCUI0333E An unexpected error occurred. The task for schedule schedule name could not be paused.	1348
BPCUI0334E An unexpected error occurred. The task for schedule schedule name could not be resumed.	1348
BPCUI0335E The volumes cannot be converted or moved because the target pools do not have sufficient available space.	1348
BPCUI0336I The ability to provision with block storage devices is only available with the advanced license.	1349
BPCUI0338E Insufficient user privileges to service the REST request.	1349
BPCUI0339E An unexpected error occurred while authorizing the REST request.	1349
BPCUI0340I A task was successfully renamed.	1349
BPCUI0341E The task could not be renamed.	1349
BPCUI0342E The task could not be renamed because the specified name already exists.	1349
BPCUI0343I Performance monitoring is unavailable for resource resource name because the resource was not probed.	1350
BPCUI0344W The following service classes allow provisioning only from the capacity pool capacity pool: service classes. If you delete this capacity pool, the service classes will allow provisioning from any available storage.	1350
BPCUI0346I The Storage Resource agent that is deployed on the server cannot be uninstalled.	1350
BPCUI0347I All servers were removed except for the product server. Entries for the product server resources might still be displayed in the GUI until all the associated removals are complete.	1350
BPCUI0348W You cannot provision volumes because at least one of the selected hosts was not found in the database. Ensure that all hosts that are selected for provisioning are being monitored.	1350
BPCUI0349W You cannot provision volumes because not all of the selected hosts appear to have Fibre Channel connectivity.	1351
BPCUI0350W You cannot provision volumes because the hypervisors that host the virtual machines use different operating systems.	1351
BPCUI0351W You cannot provision volumes because there is no Fibre Channel host port information for at least one hypervisor.	1351
BPCUI0352W You cannot provision volumes because not all of the hypervisors that host the virtual machines appear to have Fibre Channel connectivity.	1351
BPCUI0355W You cannot provision volumes because no block-storage service class exists.	1352

BPCUI0356W You cannot provision shares because no file-storage service class exists.	1352
BPCUI0357W You cannot provision volumes because you do not have permission to provision by using any block-storage service class.	1352
BPCUI0358W You cannot provision shares because you do not have permission to provision by using any file-storage service class.	1352
BPCUI0359E The credentials for the servers were not updated.	1352
BPCUI0360W The credentials for successfulUpdates of attemptedUpdates of the selected servers were updated.	1353
BPCUI0361W The credentials for the selected server was updated, however warnings did occur.	1353
BPCUI0362W The credentials for successfulUpdates selected servers were updated, however warnings did occur.	1353
BPCUI0363E Cannot connect to the SNMP data source IP_Address.	1353
BPCUI0364I The performance monitor schedule was updated for deviceName.	1353
BPCUI0366W The server serverName was not updated because it does not support the action.	1353
BPCUI0367W You cannot provision volumes to virtual machines with NPIV ports and virtual machines without NPIV ports at the same time. To provision volumes to virtual machines, ensure that you select either only virtual machines with NPIV ports or only virtual machines without NPIV ports.	1354
BPCUI0368W You cannot provision volumes because none of the selected hosts appear to have Fibre Channel connectivity and the automatic zoning option is enabled. Disable the automatic zoning option in your zoning policy.	1354
BPCUI0369W You cannot provision volumes because none of the hypervisors that manage the selected virtual machines appear to have Fibre Channel connectivity and the automatic zoning option is enabled. Disable the automatic zoning option in your zoning policy.	1354
BPCUI0370E The display name displayName is already assigned to resource resource Name.	1354
BPCUI0372I The selected hosts do not appear to have Fibre Channel connectivity. In the resulting provisioning task, ensure that the recommended storage system is connected to the hosts before you run the task. Also, be aware that all fabric-related options will be ignored.	1355
BPCUI0373I Volumes are assigned to the hypervisors that host virtual machines. Volumes are not assigned directly to virtual machines that do not have NPIV ports. None of the hypervisors that manage the virtual machines appear to have Fibre Channel connectivity. In the resulting provisioning task, ensure that the recommended storage system is connected to the hypervisors before you run the task. Also, be aware that all fabric-related options will be ignored.	1355
BPCUI0374E Schedule is not enabled for the resource resource.	1355
BPCUI0375E Performance data is not available.	1356
BPCUI0376E Invalid number of days to keep configuration history. The number should be between minimum value and maximum value.	1356
BPCUI0377E Invalid number of days to keep data for removed resources. The number should be between minimum value and maximum value.	1356
BPCUI0378E Invalid number of days to keep sample performance data. The number should be between minimum value and maximum value.	1356
BPCUI0379E Invalid number of days to keep hourly performance data. The number should be between minimum value and maximum value.	1356
BPCUI0380E Invalid number of days to keep daily performance data. The number should be between minimum value and maximum value.	1357
BPCUI0381E Failed to update the performance data retention settings.	1357
BPCUI0382E Performance monitoring is unavailable for resource resource name.	1357
BPCUI0383E Failed to update the history retention settings.	1357
BPCUI0384E Failed to retrieve the history retention settings.	1357
BPCUI0385E Invalid number of runs to keep log files for each schedule. The number should be between minimum value and maximum value.	1358
BPCUI0386E A job cannot be run for resource resourceName because there is a job already running for the resource. Wait for the job to finish and try again.	1358
BPCUI0387I The selected resources support different performance monitor intervals. If you select multiple resources, intervals that are common to all resources are displayed in the interval list.	1358
BPCUI0388E The probe schedule cannot be created for resource {0} because not all the information was provided. If you are configuring a probe for a resource for the first time, you must enter values for the probe status, time, and frequency fields.	1358
BPCUI0389E The performance monitor schedule cannot be created because not all the information was provided. If you are configuring a performance monitor for a resource for the first time, you must enter values for the performance monitor status and interval fields.	1359
BPCUI0390I The service logs were successfully created.	1359
BPCUI0391I The connection test to data source data source was successful. A probe is running. The health status is unknown until the probe is finished.	1359
BPCUI0392I The connection test to the data source data source was successful.	1359
BPCUI0393E The user user_name does not have sufficient privileges to deploy the vSphere Web Client extensionr.	1360
BPCUI0394E The user user_name does not have permission to log in to the vCenter Server system.	1360
BPCUI0395E This version of the vCenter Server server_name does not support the deployment of the vSphere Web Client extension for the product.	1360
BPCUI0396E The user user_ID does not have the required role. The role associated with this user must be Administrator, Monitor, or External Application.	1360
BPCUI0397E The vCenter Server user name or password is invalid.	1360
BPCUI0398E The user name or password is invalid.	1361
BPCUI0399I The server was started.	1361
BPCUI0400E Failed to retrieve the system management information from the Data server.	1361
BPCUI0402E Failed to retrieve the server status of the Data server.	1361
BPCUI0403E The SMI-S provider service is not available.	1361
BPCUI0404E An error occurred while updating the trace log configuration file. The original file file was deleted and could not be restored. A backup of this file may be available at backup file.	1361
BPCUI0405E Failed to set the trace settings from the Data server.	1362
BPCUI0406E Cannot start the server. The start script reported the following error: error	1362
BPCUI0407E Cannot start the server. Unable to locate the start script path to script.	1362
BPCUI0408E Cannot start the server. Unable to execute the start script path to script.	1362
BPCUI0409W The server is taking a long time to start. If the server status continues to show an error status after a reasonable interval, try to start the server again.	1363
BPCUI0410E Cannot stop the server. The stop script reported the following error: error	1363
BPCUI0411W The server is taking a long time to stop. If the server status continues to show that it is still running try to stop the server again after a reasonable interval.	1363
BPCUI0412E Cannot stop the server. Unable to locate the stop script path to script.	1363
BPCUI0413E Cannot stop the server. Unable to execute the stop script path to script.	1363
BPCUI0414W It is taking a long time for the services to start. If the server status continues to show an error status after a reasonable interval, try to start the services again. If the problem persists then restart the server.	1364
BPCUI0415E Failed to start the service service name.	1364
BPCUI0416I The server was stopped.	1364
BPCUI0417I The services of the server were started.	1364
BPCUI0418E The action cannot be completed because the data source that is managing this resource cannot be reached.	1364

BPCUI0419E A Storage Resource agent is already deployed for this server and has a status of Pending deployment or Failed deployment. Use the Servers page to resolve the deployment errors or modify the deployment schedule.	1365
BPCUI0420E A file access error occurred when the system attempted to back up or modify the tracing configuration file configuration file.	1365
BPCUI0421E There is a log collection operation already running. A new one cannot be submitted until the current one completes.	1365
BPCUI0422E Cannot start the log collecting job. Unable to locate the required script path to script.	1365
BPCUI0423E Cannot start the log collecting job. Unable to run the log collection script path_to_script.	1365
BPCUI0424E Storage cannot be provisioned from capacity pool capacity pool using service class service class for the following reason:	1366
BPCUI0425W The task task name cannot be scheduled because it is already running.	1366
BPCUI0426E Storage cannot be provisioned by using service class service class for the following reason:	1366
BPCUI0427W The selected group action is complete for all tasks, but warnings were reported.	1366
BPCUI0428I The selected group action is complete for all tasks. Some informational messages were returned.	1366
BPCUI0429E The validation process cannot contact the server. The server might be down or unreachable due to network problems.	1367
BPCUI0430I Some tasks were not deleted because they were already run.	1367
BPCUI0431E Failed to retrieve the list of managed devices.	1367
BPCUI0432E Failed to retrieve the performance monitoring granularity from the Device server. Check the connection to the Device server and retry the operation.	1367
BPCUI0433E OS type osType specified on line line of file file is not valid.	1367
BPCUI0434E Data source data_Source_Key could not be found.	1368
BPCUI0435E Required host name or IP address and OS type were not specified on line line of file file.	1368
BPCUI0436E The alert notification settings cannot be displayed.	1368
BPCUI0437E The alert notification settings cannot be saved.	1368
BPCUI0438E File file does not exist or is empty.	1369
BPCUI0439E The file file could not be uploaded.	1369
BPCUI0440E The text location specified on line line of file file has invalid character(s): characters	1369
BPCUI0441E The alert definitions cannot be displayed.	1369
BPCUI0442E The alert definitions cannot be saved.	1369
BPCUI0443E Select at least one managed server that is deployed for which alert notification settings need to be displayed.	1370
BPCUI0444E Select at least one managed server that is deployed for which alert definitions need to be displayed.	1370
BPCUI0445W The discovery job completed with errors. Some available devices were not discovered.	1370
BPCUI0446E Unable to test the connection to the device because the request was not processed by the data collector.	1370
BPCUI0447E Select at least one managed storage subsystem for which alert notification settings need to be displayed.	1370
BPCUI0448E Select at least one managed storage subsystem for which alert definitions need to be displayed.	1371
BPCUI0449E The user does not have the required authority to complete the task or command.	1371
BPCUI0451E One or more applications from provided list: names do not exist.	1371
BPCUI0452E entity name is not supporting data collection actions.	1371
BPCUI0453E One or more departments from provided list: names do not exist.	1371
BPCUI0455I No performance data is available for the selected resources.	1372
BPCUI0456E You cannot complete the action because the service is temporarily unavailable.	1372
BPCUI0457W The applications listOfApplications cannot be deleted because they contain subcomponents subcomponent, which cannot be moved up a level in the applications hierarchy due to name conflicts with existing applications in that higher level.	1372
BPCUI0458W The departments listOfDepartments cannot be deleted because they contain subdepartments or applications subdepartment, which cannot be moved up a level in the departments hierarchy due to name conflicts with departments in that higher level.	1372
BPCUI0459W The selected subcomponents cannot be removed from the application because they cannot be moved up a level in the application hierarchy due to name conflicts with the existing applications or subcomponents at the higher level.	1373
BPCUI0460W The selected applications or subdepartments cannot be removed from the department because they cannot be moved up a level in the department hierarchy due to name conflicts with the existing applications or subdepartments at the higher level.	1373
BPCUI0461W There are no task details to display. The analysis-execution task could not be run.	1373
BPCUI0462E Failed to add the device because the data collector is not responding.	1373
BPCUI0463E The discovery failed because the data collector is not responding.	1373
BPCUI0464E The connection test failed because the data collector is not responding.	1374
BPCUI0465E The requested action failed because the data collector is not responding.	1374
BPCUI0466I The servers were created.	1374
BPCUI0467W successCount of totalCount servers were created.	1374
BPCUI0468E The creation of the servers failed.	1374
BPCUI0469E Schedule job does not exist for entity name.	1374
BPCUI0470E Invalid file file size of size GB. Maximum allowed file size is max size GB.	1375
BPCUI0471E Failed to set the trace settings from the Alert server.	1375
BPCUI0472E Failed to retrieve the system management information from the Alert server.	1375
BPCUI0474E Failed to retrieve the server status of the Alert server.	1375
BPCUI0475I The volumes have been excluded from the reclamation analysis.	1375
BPCUI0476I The volumes will be included in future analyses to reclaim storage.	1376
BPCUI0477E An unexpected error occurred when modifying the optimization characteristics of the volumes.	1376
BPCUI0478E The scheduled agent upgrade time is in the past.	1376
BPCUI0479E The object storage credentials are incorrect. Enter the correct credentials. Alternatively, clear the object credentials check box and do not specify the authentication credentials for object storage now. You can use the Modify Connection action to add the object storage later.	1376
BPCUI0480E An object storage request failed on the GPFS cluster.	1377
BPCUI0481W No resources were removed.	1377
BPCUI0482E No resources were updated.	1377
BPCUI0483E The connection information cannot be updated because it points to another device.	1377
BPCUI0484I The connection information for device name was updated.	1377
BPCUI0485E The connection information cannot be updated.	1377
BPCUI0486E Cannot query the object service for information about accounts and containers as the specified user does not have admin privileges.	1378

BPCUI0487I The connection information of the selected device was successfully updated. Other devices were detected as being managed by the same data source. Would you like to update the connection information of all of them?	1378
BPCUI0488I The connection information of all devices connecting through this data source was updated.	1378
BPCUI0489W Some of the devices connecting through this data source failed to be updated.	1378
BPCUI0490I The vCenter vCenter Server was removed.	1379
BPCUI0491E The vCenter vCenter Server was not found in the database.	1379
BPCUI0492E The selected vCenter Servers were not found in the database.	1379
BPCUI0493I The vCenter vCenter Server and all number of monitored hypervisors hypervisors monitored by it were successfully removed.	1379
BPCUI0494I The number of vCenters selected vCenter Servers and all number of monitored hypervisors hypervisors monitored by them were successfully removed.	1379
BPCUI0495W Only number of removed vCenters of number of selected vCenters of the selected vCenter Servers and number of removed monitored hypervisors of number of monitored hypervisors of the hypervisors monitored by them were successfully removed.	1380
BPCUI0496I The following fabrics were detected as being managed by the same data source: comma separated fabrics list. This action applies to all fabrics that are managed by the current data source. Would you like to update the connection information of all of them?	1380
BPCUI0497E The following fabrics cannot be monitored through the SMI agent: comma separated fabrics list. The data source connection information will not be updated.	1380
BPCUI0498E The fabric cannot cannot be monitored through the SMI agent.	1380
BPCUI0499I Other switches were detected as being managed by the same data source. This action applies to all switches that are managed by the current data source. Would you like to update the connection information of all of them?	1381
BPCUI0500E One or more switches cannot be monitored through the SMI agent. The data source connection information will not be updated.	1381
BPCUI0501E The information cannot be displayed. Log out of the GUI, log in, and try the action again.	1381
BPCUI0502E The device is already managed by this data source. The data source connection information will not be updated.	1381
BPCUI0503I The connection information of the selected switches was updated.	1382
BPCUI0504I The detected versions of the resources discovered on the data source data_Source_Address are unsupported.	1382
BPCUI0505E The resource does not have a connection configured.	1382
BPCUI0506E Cannot connect to the Alert server.	1382
BPCUI0507E The version of the tpc_server IBM Spectrum Control Server is not supported.	1382
BPCUI0508E Cannot connect to the rollup server rollup_server on port host_port.	1383
BPCUI0509E Cannot authenticate with the rollup server using the provided credentials.	1383
BPCUI0510E You entered an invalid time range. The start date and time must be before the end date and time.	1383
BPCUI0511E The following alert name(s) are not unique: names.	1383
BPCUI0512E Custom alerts already exist for other resources with the following alert name(s): names.	1383
BPCUI0513E Unable to connect from rollup server rollup_server to the repository database.	1384
BPCUI0514E The specified secondary server secondary_server is the primary server.	1384
BPCUI0515E The duration of the automated probe run window must be at least minimum_hours hours.	1384
BPCUI0516W The selected subgroups cannot be removed from the general group because they cannot be moved up a level in the groups hierarchy due to name conflicts with the general groups at the higher level.	1384
BPCUI0519E Authorization has failed because the private key is not valid for the user name that you have specified.	1384
BPCUI0520E The IP address ip_address for the FlashSystem storage system is not the management IP address.	1385
BPCUI0521E The configuration for the report can't be saved.	1385
BPCUI0522E Failed to delete a report configuration.	1385
BPCUI0523E Alerts cannot be defined for this storage system.	1385
BPCUI0524E The changes to the report configuration can't be saved.	1385
BPCUI0525E The configuration for the report can't be saved because the report title isn't unique.	1386
BPCUI0527E The action cannot be completed because of an invalid request.	1386
BPCUI0528E The action cannot be completed because of an invalid file upload request.	1386
BPCUI0526I The connection test to data source data source was successful. A probe is running.	1386
BPCUI0529I The data source data_Source_Address is already being managed as a data source for monitoring. No new resources were detected.	1386
BPCUI0530I The data source data_Source_Address is already being managed as a data source for monitoring. The following new resources were detected:	1387
BPCUI0531E The action cannot be completed because LDAP registry file failed to upload.	1387
BPCUI0532E The action failed because of a missing resource.	1387
BPCUI0533E The LDAP configuration test failed.	1387
BPCUI0534E There was an error executing the collect log process. If this problem persists, you can try collecting and uploading the service logs manually. Learn More.	1388
BPCUI0535E An FTP connection can not be established. If your organization requires the use of a proxy server, consult the following documentation: Troubleshooting FTP Transfers.	1388
BPCUI0536E The support data collection failed due to an invalid PMR number format.	1388
BPCUI0537E The support package could not be created because file system permissions prevent the creation of temporary files.	1388
BPCUI0538E The support data collection completed creating a support package, but the package could not be uploaded to IBM.	1388
BPCUI0539E The support data collection failed with an internal error	1389
BPCUI0540E The support data collection failed due to an invalid email address format.	1389
BPCUI0541E The specified SMI agent was not found. Make sure that the protocol, SMI agent host name or IP address, and port are specified correctly and that the SMI agent is properly configured at that location.	1389
BPCUI0542E A connection was not established. Make sure that the protocol, SMI agent host name or IP address, and port are specified correctly.	1389
BPCUI0543E The authentication to the SMI agent failed.	1389
BPCUI0544E There is a pending delete in process for this SMI agent.	1390
BPCUI0545E The SMI agent service is not available.	1390
BPCUI0546E The action cannot be completed because the LDAP registry file could not be updated.	1390
BPCUI0547E Connection failed. The server might be down or unreachable due to network problems.	1390
BPCUI0548E The add SSL certificate action failed.	1390
BPCUI0549E The add SSL certificate action failed because of a wrong password.	1390
BPCUI0550E The specified storage resource is not valid for the REST API service request.	1391
BPCUI0551E The file cannot be used because it is not a valid SSL certificate. Select a valid certificate file and try again.	1391

BPCUI0554E The SSL certificate download process failed.	1391
BPCUI0555E The test connection to the LDAP server failed. Verify that your XML file contains the correct syntax and values and that the LDAP server is running.	1391
BPCUI0556E An unexpected error occurred creating or updating a support ticket.	1391
BPCUI0557E An invalid request was made when creating or updating a support ticket.	1392
BPCUI0558E This tier name is already in use. Enter a different name.	1392
BPCUI0559E The custom dashboard was removed by another user. Cancel the action and refresh the page manually.	1392
BPCUI0600W Can't save the scheduling information for the report because the Data server is offline.	1392
BPCUI0601I The resource does not have a connection configured. To add a connection to the resource, click Add Switch or Add Fabric.	1392
BPCUI0602E The osAuthentication script does not start. The script reported the following error: script_error.	1393
BPCUI0603E The connection test to data source data source was not successful.	1393
BPCUI0604E Can't stop data collection for entity name.	1393
BPCUI0605E Can't restart data collection for entity name.	1393
BPCUI0606E The action cannot be completed because there was a failure to create or write into the pending configuration file.	1394
BPCUI0607E The action cannot be completed because there was a failure to read the pending LDAP registry file.	1394
BPCUI0608E The action cannot be completed because there was a failure to get the list of LDAP groups.	1394
BPCUI0609E The Local OS authentication configuration test failed.	1394
BPCUI0610E Failed to update modified username for IBMid unique ID.	1394
BPCUI0611E Failed to update IBMid unique ID for IBMid username.	1395
BPCUI0612E The action was not performed due to invalid device credentials for entity name.	1395
BPCUI0613E A switch with this host name or IP address is already being monitored.	1395
BTACD - Database verifier messages for SAN database service	1395
BTACD0010I IBM Spectrum Control (Configuration Service) initialized successfully.	1395
BTACD0011E An error occurred while reading properties from file file name	1395
BTACD0012I Properties were successfully read from file file name	1396
BTACD0013E Exception occurred while saving the properties file file name.	1396
BTACD0014E An error was received while attempting to remove a callback key.	1396
BTACD0015E An error was received while attempting to get the database driver driver name.	1396
BTACD0016E An error was received while decrypting the database password.	1397
BTACD0017E There was an error in setting the WAS Admin password.	1397
BTACE - SAN event services messages	1397
BTACE0030I IBM Spectrum Control (Event Service) initialized successfully.	1397
BTACE0031I A SAN Event message was received from: publisher name.	1397
BTACE0032E Missing SNMP destination address. SNMP trap was not sent.	1398
BTACE0033E An error occurred when trying to send an event to the Tivoli Enterprise Console server.	1398
BTACE0034E The Tivoli Enterprise Console server location has not been specified.	1398
BTACE0035E JMSEException: exception	1398
BTACE0037E An error occurred during the process of forwarding an SNMP trap or sending a Tivoli Enterprise Console event.	1398
BTACE0039E Unable to save the SnmpAddress entry to the properties file properties file.	1399
BTACE0040E Unable to remove the SnmpAddress entry from the properties file properties file.	1399
BTACE0041E Unable to save the SAN Domain ID to the properties file properties file.	1399
BTACE0042E Unable to read the SNMP and Tivoli Enterprise Console server address entries from the properties file properties file.	1399
BTACE0507E Messaging Service is not running and cannot be used to subscribe or publish events.	1399
BTACE0508E A proxy to the Messaging Service could not be obtained.	1400
BTACE0509E The proxy to the Messaging Service might not be valid. Cannot publish or subscribe to events.	1400
BTACE0510E Unable to create topic topic name.	1400
BTACE0511E Unable to create a publisher for topic topic name.	1400
BTACE0512E Unable to create a subscriber for topic topic name.	1400
BTACE0513E Unable to create a message.	1401
BTACE0514E Service service name failed to subscribe to topic topic name.	1401
BTACE0515E Publish failed for topic topic name.	1401
BTACS - Service manager messages	1401
BTACS0001I IBM Spectrum Control command line interface initialized successfully.	1402
BTACS0002E The -url requires a host:port argument.	1402
BTACS0003E Required parameters are not present.	1402
BTACS0002I The command line interface is binding to the service.	1403
BTACS0003I The bind was successful.	1403
BTACS0004E An exception occurred while invoking the service: method name.	1403
BTACS0015E Caught exception: value.	1403
BTACS0005E Could not deserialize exception: value.	1403
BTACS0006E Fault detail: value.	1403
BTACS0014E An undeclared exception was encountered: value.	1404
BTACS0007E Unrecognized command for service: service name.	1404
BTACS0005I Deployed service service name: class=value, scope=value, autostart=value, static=value, order=value.	1404
BTACS0006I Undeploying service: service name	1404
BTACS0007I Undeployed service: service name	1404
BTACS0008E Error undeploying service value : value	1404
BTACS0009E Error starting the service name service.	1405
BTACS0004I Started service service name.	1405
BTACS0011E Service service name was not deployed.	1405
BTACS0008I Starting service service name (timeout number seconds)	1405
BTACS0013E Service service name did not start, interrupting the startup thread.	1405
BTACS0009I The service service name was stopped.	1406

BTACS0010E An error was encountered while stopping service service name.	1406
BTACS0012E Exception was received while stopping service value: value.	1406
BTACS0016E Service service name is not available.	1406
BTACS0010I Stopping service service name (timeout number seconds).	1406
BTACS0017E Service service name did not stop, interrupting the startup thread.	1407
BTACS0018E A problem was encountered while getting class definition: class definition name.	1407
BTACS0019E An interface value in service service name is being ignored.	1407
BTACS0020E Error starting service value: value.	1407
BTACS0021E Exception getting status from service value: value.	1407
BTACS0011I Interrupting monitor thread and waiting for it to exit.	1408
BTACS0012I Service Manager shutting down.	1408
BTACS0013I Monitoring services (monitor interval is number seconds).	1408
BTACS0014I Service service name has value.	1408
BTACS0015I The Service Manager monitor process is exiting.	1408
BTACS0016I There are no services to autostart.	1408
BTACS0017I All autostart services have started.	1408
BTACS0018I All services are shutting down.	1409
BTACS0019I All services have been shut down	1409
BTACS0020I Deleted file file name	1409
BTACS0021I Unable to delete file file name	1409
BTACS0022I Starting autostart services.	1409
BTACS0023I An error occurred while starting the service name service.	1409
BTACS0024I The properties from file file name were successfully read.	1409
BTACS0025E An error occurred while reading properties from file file name.	1409
BTACS0026E Login failed: Unknown user name or bad password.	1410
BTACS0027E Command failed: Failed to connect.	1410
BTACS0028E Command failed.	1410
BTACS0029E Failed to connect to <hostname>.	1410
BTACS0030E Failed to authenticate host <hostname>.	1411
BTACS0031I The server is not registered with the Agent Manager.	1411
BTACS0032I Registering with the Agent Manager at host name:port.	1411
BTACS0033I The server is renewing credentials with the Agent Manager at host name:port.	1411
BTACS0034I The server credentials are current. Agent Manager at host name:port.	1411
BTACS0035E The server failed to register with the Agent Manager at host name:port.	1411
BTACS0036W The server failed to register with the Agent Manager. It will retry in delay seconds.	1412
BTACS0037I The server successfully registered with the Agent Manager.	1412
BTACS0038I The server successfully renewed the credentials with the Agent Manager.	1412
BTACS0039W The server failed to renew the credentials with the Agent Manager.	1412
BTACS0040E This command requires additional arguments.	1412
BTACS0041E The command line is not available for service: service name.	1412
BTACS0042E Invalid command: CLI command	1412
BTACS0043E Failed to authenticate with host hostUrl. Invalid host authentication password.	1412
BTACS0044E The server failed to register with the Agent Manager: Incorrect agent registration password.	1413
BTACS0045I SERVICE MANAGER COMMANDS	1413
BTACS0046I Returns the status of the services.	1413
BTACS0047I Service.functionName performed by user at location. Input parameters: input parameters, output parameters: output parameters	1413
BTACS0048W Unauthorized request by user at location to perform service.functionName.	1413
BTACS0049W Not licensed to perform service.functionName request by user at location.	1414
BTACS0050I Waiting for Common Agent services.	1414
BTACS0051I The Common Agent services have started successfully.	1414
BTACS0052W Failed to create the Common Agent service filter. The Fabric agent will start without waiting for the required Common Agent services.	1414
BTACS0053I Agent startup is already in progress.	1414
BTACS0054I Invalid Server ID has been provided to update Server Job status.	1414
BTACS0055I Agent Manager Registration is set to NO. The server will not register with an AgentManager.	1414
BTACS0056I Agent Manager Registration is set to YES. The server will register with the AgentManager.	1415
BTACS0057W Error in configuration parameter AgentManager.Registration - default value will be used.	1415
BTACS0000I Starting Control Process: value, Device Server RUN ID=value, Job ID=value.	1415
BTACS0001I IBM Spectrum Control command line interface initialized successfully.	1402
BTACS0002I The command line interface is binding to the service.	1403
BTACS0003I The bind was successful.	1403
BTADS/BTAFM/BTAVM/HWN - Job logging messages	1415
BTADS0000I Starting Discover Process value , with Device Server RUN ID value , and Job ID value .	1431
BTADS0001I Discover Process with Device Server run ID value and job ID value is complete.	1432
BTADS0002I Starting Child Discover Process value with Job ID= value .	1432
BTADS0003I The Child Discover Process with Job ID value has completed with Status= value and Return Code= value .	1432
BTADS0004W The child discovery request with job ID job_id completed with status status_number and return code value.	1432
BTADS0005E The child discovery request with job ID job_id completed with status status_number and return code value.	1432
BTADS0010I Invoking outband scanner value on agent value .	1433
BTADS0011I Outband scanner value on agent value completed successfully.	1433
BTADS0012E Outband Scanner value on agent value failed with return code value .	1433
BTADS0019E An outband scanner failed to capture the scan data.	1433
BTADS0020I Processing value data from agent value .	1433

BTADS0021W Warning encountered while parsing Fabric XML for job: RUN ID= value , and Job ID= value . value .	1434
BTADS0022E Exception encountered while parsing Fabric XML for job: RUN ID= value , and Job ID= value . value .	1434
BTADS0023E Fatal error encountered while parsing Fabric XML for job: RUN ID= value , and Job ID= value . value .	1434
BTADS0024E Error encountered processing scanner value data from agent value . value .	1434
BTADS0025I Running job to discover SMI-S providers through Service Location Protocol: RUN ID= value , Job ID= value .	1435
BTADS0026I Service Location Protocol has found value SMI-S providers.	1435
BTADS0027E Error encountered by a Service Location Protocol job: RUN ID= value , and Job ID= value . value .	1435
BTADS0028W The Device Server Job with RUN ID=: value , Job ID= value , Discover Request= value has been cancelled since it is long running.	1435
BTADS0029I Scanner value data from agent value has not changed since last scan.	1435
BTADS0030I Invoking inband Scanner value on agent value .	1436
BTADS0031I Inband Scanner value on Agent value completed successfully.	1436
BTADS0032E Inband Scanner value failed on agent value with Return Code value .	1436
BTADS0033E Error invoking value on host value .	1436
BTADS0034E Fatal error encountered while persisting the data for job: RUN ID= value , and Job ID= value . value .	1436
BTADS0035E The execution of the job failed with: value .	1437
BTADS0036I Found SNMP Target at value .	1437
BTADS0037E Found SNMP Target at value but it is not persisted in the database. Will NOT perform discovery of information using the address.	1437
BTADS0038I Starting scan of SNMP agents from value to value .	1437
BTADS0039I Starting probe of detected agents.	1437
BTADS0040I Processing of Scanner value data from Agent value completed successfully.	1438
BTADS0041I Discover Process with Device Server RUN ID value and Job ID value completed successfully.	1438
BTADS0042E Discover Process with Device Server RUN ID value and Job ID value failed with return code value .	1438
BTADS0043I Invoking value scanner value on target value .	1438
BTADS0044I value scanner value on target value completed succesfully.	1438
BTADS0045E value Scanner value on target value failed with return code value .	1439
BTADS0046I Processing value data from agent value .	1439
BTADS0047W The value parser encountered a warning while parsing XML for job with RUN ID= value , and Job ID value . The return code from the parser job is value .	1439
BTADS0048E The value parser encountered an exception while parsing XML from job with RUN ID= value , and Job ID= value .The return code from the parser is value .	1439
BTADS0049E The value parser for Device Server job with RUN ID= value , and Job ID= value failed. The return code from the parser is value .	1440
BTADS0050I Service Location Protocol has found SMI-S provider, value , at address value .	1440
BTADS0051I Service Location Protocol has found SMI-S provider, value , at address value , which requires security information to be configured.	1440
BTADS0052W Warning encountered while parsing value data from agent value. value.	1440
BTADS0053E Exception encountered while parsing value data from agent value. value.	1440
BTADS0054E Fatal error encountered while parsing value data from agent value. value.	1441
BTADS0055E Outband Scanner value on agent value encountered the presence of a McData i10k. These devices do not report correctly via SNMP and can only be used with SMI-S provider.	1441
BTADS0056E Errors in Topology XML generator.	1441
BTADS0057E Errors occurred while resolving InterconnectElement and Port relationship.	1441
BTADS0058E Errors in creating an entity.	1442
BTADS0059E The outband agent target address IP address is not a Cisco device or is invalid.	1442
BTADS0060E Outband Scanner value is not responding.	1442
BTADS0062E Encountered SQL error value while persisting some data.	1442
BTADS0063E The execution of the PM BSP invocation failed with: value .	1443
BTADS0063W The performance data collection for the current device is not enabled.	1443
BTADS0064I Starting scan of Storage Subsystems from value to value .	1443
BTADS0065I Outband and inband agents for fabric(s) specified in probe are value	1443
BTADS0066I Could not find scanners for agent value	1443
BTADS0067I Agent value is configured for no SAN calls and so no scanners will be invoked for this particular agent	1444
BTADS0068I Could not retrieve connection information for agent value. Will not be able to invoke scanners for this particular agent	1444
BTADS0069I Added inband scanner job with id value discover request value for agent value.	1444
BTADS0070I Agent value has not discovered any fabrics and will not be used during the probe.	1444
BTADS0071I Invoked inband Scanner value on agent value .	1444
BTADS0072I Successfully received response from agent for job value with request id value .	1445
BTADS0073E Received error response from agent for job value with request id value. Return code is value.	1445
BTADS0074E IP Scan Discovery was canceled due to a hung socket/thread detected. Partial result of the scan will be persisted.	1445
BTADS0075E IP Scan Discovery was canceled due to a hung socket/thread detected.	1445
BTADS0076I IP Scan Discovery has started for DS, XIV, and IBM SONAS subsystems.	1445
BTADS0077I Scanning value out of value IP addresses.	1446
BTADS0078I IP Scan Discovery has started for SVC subsystems.	1446
BTADS0079I IP Scan Discovery for DS and XIV was done.	1446
BTADS0080I IP Scan Discovery for SVC was done	1446
BTADS0081I Inband Scanner value for agent address value is not required for probing switches and will not be used.	1446
BTADS0082W A first run of a switch probe failed. Additional agents will be used.	1446
BTADS0083I The available agents provide a subset of possible features for the probed switch: value	1447
BTADS0084I There are no limitations for probing switch value based on the mix of agents that are configured.	1447
BTADS0085W A problem was encountered when agent assignments were being determined for the probe. All available agents will be used to collect information about the switch.	1447
BTADS0086I The following storage systems were discovered value	1447
BTADS0087I IP Scan Discovery did not find any DS8000, SVC, XIV, and IBM SONAS storage systems in the given IP range.	1447
BTADS0088I IP Scan Discovery finished.	1448

BTADS0089E The Device server is not registered with agent manager. Scanners cannot be used for agent value. value. _____	1448
BTADS0090E There are no agents currently available to probe switch value. _____	1448
BTADS0091I Inband Scanner value for agent address value is currently not running and will not be used. _____	1448
BTADS0092I Inband Scanner value for agent address value is currently disabled from performing fabric functions and will not be used. _____	1449
BTADS0093I Inband Scanner value for agent address value is currently not reachable and will not be used. _____	1449
BTADS0094W The probe for switch value has some limitations. _____	1449
BTADS0095W For switch value some information will not be collected. _____	1449
BTADS0096I The probe limitation can be overcome by configuring an SMI agent to manage fabric value. _____	1449
BTADS0097I The probe limitation can be overcome by configuring SNMP agents to manage switches in fabric value. _____	1450
BTADS0098I The probe limitation can be overcome by configuring a Storage Resource agent to manage fabric value. _____	1450
BTADS0099W The following WWN is not recognized as belonging to a known vendor: value. _____	1450
BTADS0100W Invalid relationships between switches and fabrics were identified. If possible, these relationships will be fixed automatically for the following switches: value. _____	1450
BTADS0101W The discover process that has the Device server run ID value and job ID value completed with one or more warnings. _____	1451
BTADS0102E The probe with the run ID value completed with errors. _____	1451
BTADS0103E No data source is available to probe switch switch_name. _____	1451
BTADS0104E A timeout occurred while processing the request. Try the request again. _____	1451
BTADS0105E A response from the data collector was not received within the specified time. _____	1452
BTADS0106E The requested action on agent agent_name did not complete because the data collector stopped or is not responding. The request failed with error code error_code. _____	1452
BTADS0107W Outband Scanner outband_scanner_name on agent agent_name failed because of another transaction in progress on the switch. _____	1452
BTADS0108E Outband Scanner outband_scanner_name on agent agent_name failed because unexpected data was returned by the switch. Check the trace file for more details. _____	1452
BTADS0109I Outband Scanner outband_scanner_name on agent agent_name did not collect zoning data. _____	1452
BTADS0110I Outband Scanner outband_scanner_name on agent agent_name did not pass write capabilities check. _____	1453
BTADS0111E The probe was unable to collect some details of the switch. _____	1453
BTADS0112E Error encountered while persisting some data. value _____	1453
BTADS0113E Error encountered while processing a probe job. value _____	1453
BTADS0114E The information cannot be saved to the database repository. _____	1454
BTADS0115E The probe failed when collecting information about the resource. The data collector returned the following error status: value. _____	1454
BTAfM0000I Operation op_name processed successfully. _____	1454
BTAfM0100I Initializing Collection. _____	1454
BTAfM0110I Querying the SMI-S provider. _____	1454
BTAfM0113I Collecting for db_table, current_obj of num_objs. _____	1455
BTAfM0114I Probing data for switch switch_name. _____	1455
BTAfM0115I Probing data for port port_name. _____	1455
BTAfM0150I Storing Information. _____	1455
BTAfM0151I The db_table of current_obj num_objs stored. _____	1455
BTAfM0200I Traversing fabric topology. _____	1456
BTAfM0500I The IBM Spectrum Control Device Server service has started successfully. _____	1456
BTAfM0501I The IBM Spectrum Control Device Server service was shut down successfully. _____	1456
BTAfM0502I The IBM Spectrum Control Device Server service provides methods to collect, report and configure the fabric hardware. _____	1456
BTAfM0505I The delete missing function has started. _____	1456
BTAfM0506I The delete missing method was processed in milliseconds milliseconds. _____	1456
BTAfM0723W No blades were discovered for the slot slot. _____	1457
BTAfM2000W Operation op_Name partially processed. _____	1457
BTAfM2501W Unable to shut down Device Server Service smoothly. _____	1457
BTAfM4000E Operation op_Name failed. _____	1457
BTAfM4001E An internal error occurred. _____	1457
BTAfM4002E Could not get requested information due to an internal error - errorMessage _____	1458
BTAfM4100E Mandatory parameter parameter_Name is missing. _____	1458
BTAfM4101E Invalid parameter parameter_name. _____	1458
BTAfM4103E Entity entity_name was not found. _____	1458
BTAfM4104E Attribute attribute_name was not found. _____	1458
BTAfM4105E Computer computer_name was not found. _____	1459
BTAfM4106E Fabric fabric_name was not found. _____	1459
BTAfM4107E Switch switch_name was not found. _____	1459
BTAfM4108E Port port_name was not found. _____	1459
BTAfM4109E Zone set zoneset_name was not found. _____	1459
BTAfM4110E Zone zone_name was not found. _____	1459
BTAfM4111E Zone alias zone_alias_name was not found. _____	1460
BTAfM4112E Zone member zone_member_name was not found. _____	1460
BTAfM4113E Subsystem subsystem_name was not found. _____	1460
BTAfM4114E Host Bus Adapter HBA_name was not found. _____	1460
BTAfM4115E Node node_name was not found. _____	1460
BTAfM4116E Link from port from_port_name to port to_port_name was not found. _____	1460
BTAfM4117E Hub hub_name was not found. _____	1461
BTAfM4118E Router router_name was not found. _____	1461
BTAfM4119E Bridge bridge_name was not found. _____	1461
BTAfM4120E LUN LUN_name was not found. _____	1461
BTAfM4140E Agent Agent_name was not found. _____	1461
BTAfM4141E Scanner scanner_name on agent agent_name was not found. _____	1461

BTAFM4142W Agent agent_name was ignored because the switch switch_name was probed by agent agent1_name.	1462
BTAFM4150E Indexed properties property_name don't match.	1462
BTAFM4180E Agent to gather sensor and event data is not available for the switch switch_name.	1462
BTAFM4200E Credentials not found.	1462
BTAFM4300E The connection to the SMI agent for switch switch_name could not be made.	1462
BTAFM4301E The invocation of CIM method method_name failed on SMI-S provider SMI-S provider name. The return code is return_code.	1463
BTAFM4302E The invocation of CIM method method_name failed on SMI-S provider SMI-S provider name with the following exception text: exception_text.	1463
BTAFM4303E Received unexpected values from SMI-S provider SMI-S provider name .	1463
BTAFM4304E SMI agent SMI agent name can not contact switch switch_name.	1463
BTAFM4305E The CIM method method_name is not supported on the switch switch_name.	1464
BTAFM4306E Could not create connection to SMI-S provider SMI-S provider name . Reason: reason.	1464
BTAFM4307E The username user_name or password is wrong on SMI-S provider SMI-S provider name.	1464
BTAFM4308I Could not create connection to SMI-S provider SMI-S provider name . Reason: reason. An alternate SMI-S provider will be used.	1464
BTAFM4501E No agent is available to configure the zoning on the fabric with ID fabric_name.	1465
BTAFM4502E The fabric with ID fabric_name is currently locked by another client of IBM Spectrum Control.	1465
BTAFM4503E A token for fabric fabric_name has expired for client client_name.	1465
BTAFM4504E The transaction for fabric fabric_name has expired.	1465
BTAFM4505E Another transaction is in progress for fabric fabric_name.	1465
BTAFM4506E Zone set zoneset_name already exists.	1465
BTAFM4507E Zone zone_name already exists.	1466
BTAFM4508E Zone alias zone_alias_name already exists.	1466
BTAFM4509E Zone member zone_member_name already exists.	1466
BTAFM4510E Another job is in progress for fabric fabric_name.	1466
BTAFM4550E The Device Server encountered an error accessing the database.	1466
BTAFM4600E Unable to start the Device Server service.	1467
BTAFM5000E Step failed after collecting Count of collected entities entities for switch switch where entities exist. Continuing with next step.	1467
BTAFM5001E No set of fabrics or switches was defined for this probe.	1467
BTAFM5002E The SMI agents SMIURL returned an error or can no longer contact the switches.	1467
BTAFM5003E Requests to an SMI agent did not correctly collect a set of switches for fabric fabric identity.	1468
BTAFM5004E No switch retrieved from the SMI agent for fabric fabric identity.	1468
BTAFM5005E No switch found for fabric fabric identity.	1468
BTAFM5006E No switch retrieved from database.	1468
BTAFM5007E Failed to get CIM entity for fabric fabric_name.	1469
BTAFM5008E Failed to get CIM entity for switch switch_name.	1469
BTAFM5009E Failed to enumerate CIM entity Entity class name.	1469
BTAFM5010E SMI-S provider is not available.	1469
BTAFM5011E Failed to get blade for switch Switch name.	1470
BTAFM5012E Failed to get physicalpackage for blade with slot number Blade slot name.	1470
BTAFM5013E Blade serial number is NULL.	1470
BTAFM5014E Step failed after collecting Count of collected entities entities for fabric fabric where entities exist. Continuing with next step.	1470
BTAFM5015E Data source could not be retrieved from the IBM Spectrum Control database for fabric fabric where data source exists.	1471
BTAFM5016E The selected data source could not be contacted for fabric fabric where data source exists.	1471
BTAFM5017E Failed to get fabric for switch Switch name.	1471
BTAFM5018E Failed to get CIM entity for port port_name.	1471
BTAFM5019E Failed to get switch for port port_name.	1472
BTAFM5020E Failed to get blade for port port_name.	1472
BTAFM5021E Failed to get CIM entity for blade blade_name.	1472
BTAFM5022E Failed to get switch for blade blade_name.	1472
BTAFM5023E Failed to discover Fabric and Switch.	1472
BTAFM5024E The data source for switch switch_name was not retrieved from the database repository.	1473
BTAFM0600I Count of collected entities blades collected for switch switch where entities exist.	1473
BTAFM0601I Starting collection of switch blades for switch switch identifier.	1473
BTAFM0602I Collection of switch blades completed. Count of collected entities entities collected in total for switch switch identifier.	1473
BTAFM0603I Starting collection of switch fcports for switch switch identifier.	1474
BTAFM0604I Collection of switch fcports completed. count of collected entities entities collected in total for switch switch identifier.	1474
BTAFM0605I Start probing switch entities switches.	1474
BTAFM0606I Start topology probing for fabric fabric entity.	1474
BTAFM0609I Count of entities fcports collected for switch switch where entities exist.	1474
BTAFM0614I The probe task is to probe topology and zone. The probe algorithm is CIM association.	1475
BTAFM0616I The probe policy involves discovering segmented or merged fabrics.	1475
BTAFM0617I The probe policy doesn't involve discovering segmented or merged fabrics.	1475
BTAFM0618I The probe task is to probe topology. The probe algorithm is CIM association.	1475
BTAFM0620I Start zone probing for fabric fabric entity.	1475
BTAFM0621I Starting collection of zone set for switch switch entity.	1476
BTAFM0622I Starting collection of zone for switch switch entity.	1476
BTAFM0623I Starting collection of zone alias for switch switch entity.	1476
BTAFM0624I Starting collection of zone member from zone alias for switch switch entity.	1476
BTAFM0625I Starting collection of zone member and zone alias from zone for switch switch entity.	1476
BTAFM0626I Starting collection of zone member from zone for switch switch entity.	1477
BTAFM0627I Starting collection of zone set for fabric fabric entity.	1477
BTAFM0628I Count of collected entities zone sets collected.	1477
BTAFM0629I Collection of zone set completed. Count of collected entities entities collected in total for fabric fabric entity.	1477

BTAFM0630I Starting collection of zone for fabric fabric entity.	1477
BTAFM0631I Count of collected entities zones collected.	1477
BTAFM0632I Collection of zone completed. Count of collected entities entities collected in total for fabric fabric entity.	1478
BTAFM0633I Starting collection of zone alias for fabric fabric entity.	1478
BTAFM0634I Count of collected entities zone aliases collected.	1478
BTAFM0635I Collection of zone alias completed. Count of collected entities entities collected in total for fabric fabric entity.	1478
BTAFM0636I Starting collection of zone member from zone alias for fabric fabric entity.	1478
BTAFM0637I Starting collection of zone member and zone alias from zone for fabric fabric entity.	1479
BTAFM0638I Starting collection of zone member from zone for fabric fabric entity.	1479
BTAFM0639I Collection of zone member completed. Count of collected entities entities collected in total for fabric fabric entity.	1479
BTAFM0640I Zone probe will discover both active and inactive zone definitions at selected data source datasource name for zone probe.	1479
BTAFM0641I Zone probe will discover only active zone sets at data source datasource name for zone probe.	1479
BTAFM0654I The port is not switch port.	1480
BTAFM0655I The switch profile doesn't support this switch switch_name. No further process to probe this switch.	1480
BTAFM0656I Start enumerating entity of association between fabric and zone set at selected data source Url entity.	1480
BTAFM0657I Start enumerating entity of association between fabric and zone at selected data source Url entity.	1480
BTAFM0658I Start enumerating entity of association between fabric and zone alias at selected data source Url entity.	1480
BTAFM0659I Start enumerating entity of association between switch and zone set at selected data source Url entity.	1481
BTAFM0660I Start enumerating entity of association between switch and zone at selected data source Url entity.	1481
BTAFM0661I Start enumerating entity of association between switch and zone alias at selected data source Url entity.	1481
BTAFM0662I Start enumerating associations between virtual fabric and zoning entities at selected data source Url entity.	1481
BTAFM0663I Starting collection of switch control processor blades for switch switch identifier.	1481
BTAFM0664I Count of collected entities control processor blades collected for switch switch where entities exist.	1482
BTAFM0665I Collection of switch control processor blades completed. Count of collected entities entities collected in total for switch switch identifier.	1482
BTAFM0666I Checksums for the active and defined Zone Database could not be updated for fabric entity.	1482
BTAFM0667E Job id or request id is missing for a SRA job that is been processed.	1482
BTAFM0668E Command and/or job timestamp is missing for job id with request id .	1482
BTAFM0669I job id with request id was is not found. Device server may have been restarted after job was created.	1483
BTAFM0670E could not retrieve output file for job id with request id .	1483
BTAFM0671E Another probe of fabric The Name+Nameformat of the fabric is already in progress.	1483
BTAFM0672E Device server is not registered with agent manager. Will not be able to invoke scanner on host .	1483
BTAFM0673E There are no agents that are currently available to probe fabric .	1483
BTAFM0674W No fabric found for event source that is associated with switch with IP address .	1484
BTAFM0675E Unable to start parsing of SRA fabric probe data for SRA job id request id file name .	1484
BTAFM0676E Error parsing SRA fabric probe data for SRA job id request id file name .	1484
BTAFM0677E Unable to connect to SNMP port (another application may already be connected and forwarding messages).	1484
BTAFM0678I The Name of the switch switch was removed.	1484
BTAFM0679I The The Name+Nameformat of the fabric fabric was removed.	1485
BTAFM0680E The Name of the switch switch was not removed because it is not missing.	1485
BTAFM0681E The The Name+Nameformat of the fabric fabric was not removed because it is not missing.	1485
BTAFM0682E An error occurred while checking for access to the database to save new zoning information for fabric to the database.	1485
BTAFM0683E Unable to access the database to save zoning information for fabric . Another job is currently saving new zoning information to the database for the same fabric.	1485
BTAFM0684I The job is waiting to access the database to save new zoning information for fabric . Another job is currently saving zoning information to the database for the same fabric.	1486
BTAFM0685W Host/IP Address is not a switch.	1486
BTAFM0686W Switch is not a supported switch.	1486
BTAFM0687W The switch does not respond to SNMP queries.	1486
BTAFM0688W Cannot communicate with host or IP address .	1487
BTAFM0689W No ports were discovered for the switch .	1487
BTAFM0690I Collection of data from trunks is completed. Data was collected from count of collected entities trunks.	1487
BTAFM0691I Starting collection of data from trunks for switch switch identifier.	1487
BTAFM0692I Count of entities trunks collected for switch switch where entities exist.	1487
BTAFM0692E A response from the data collector was not received within the specified time.	1487
BTAFM0693E A response from the data collector was not received. The request failed with return code return_code	1488
BTAFM0694W Zoning data cannot be collected because there is a transaction in progress on the switch key	1488
BTAFM0695E The switch key is returning unexpected data.	1488
BTAFM0696E Zone set zoneset_name is already active.	1488
BTAFM0697E Zone set zoneset_name is already inactive.	1488
BTAFM0698E On the switch switch_name VSAN vsan_name was not found.	1489
BTAFM0699E The switch key did not return zoning data.	1489
BTAFM0700E Duplicate entries for the same switch: switch.	1489
BTAFM0701E Current active full zone configuration is not synchronized with the zone configuration on the switch switch_name for VSAN vsan_name .	1489
BTAFM0702E You cannot monitor Brocade Access Gateway switches without Network Advisor.	1489
BTAFM0703I Waiting for probes of other Access Gateway switches to complete.	1490
BTAFM0704W Distributing zone configuration across all the switches for VSAN vsan_name did not succeed on the switch switch_name .	1490
BTAFM0705W Zone data collection after zone changes were made failed on the switch switch_name .	1490
BTAFM0706E The fabric probe was unable to collect some details of the blades on the switches.	1490
BTAFM0707I You cannot use IBM Spectrum Control to make zoning changes for provisioning on switch switch_name.	1490
BTAFM0708E The probe was unable to collect some details of the switches.	1491
BTAFM0709I Started to process information for fabric fabric_name.	1491
BTAFM0710I Started to process information for switch switch_name.	1491

BTAFM0711I Started to process information for discovered switches.	1491
BTAFM0712I Started to process information for a switch blade.	1492
BTAFM0713I Started to process information for a switch zone set.	1492
BTAFM0714I Started to process information for switch ports.	1492
BTAFM0715E Error occurred while processing information for fabric fabric_name.	1492
BTAFM0716E Error occurred while processing information for virtual fabric virtual_fabric_name.	1492
BTAFM0717E Error occurred while processing information for switch switch_name.	1492
BTAFM0718E Error occurred while processing information for discovered switches.	1493
BTAFM0719E Error occurred while processing information for logical switches.	1493
BTAFM0720E Error occurred while processing information for active zone set active_zone_set_name.	1493
BTAFM0721E Error occurred while processing information for inactive zone set inactive_zone_set_name.	1493
BTAFM0722E Error occurred while processing information for port port_name.	1493
BTAQE1107E InbandScanHandler failed to start InbandScanner scanner name on managed host target.	1494
BTAQE1108E InbandScanHandler failed to get callback information for InbandScanner scanner name on managed host target.	1494
BTAQE1112E During an outband scan, the scanner scanner name was unable to identify the target host target.	1494
BTAQE1113E Unable to invoke an Outband scan scanner name on target target.	1494
BTAQE114E OutbandScannerHandler received invalid callback information for Outband scanner scanner name on target target.	1495
BTAQE1115E The outband scanner scanner name did not return the SAN ID on target target.	1495
BTAVM0001I The operation Name of the operation processed successfully.	1495
BTAVM0002I The Web service call Name of the operation processed successfully.	1495
BTAVM0003I Data source Name of the datasource successfully added.	1496
BTAVM0004I Data source Name of the datasource successfully deleted.	1496
BTAVM0005I Data source Name of the datasource successfully modified.	1496
BTAVM0006I Discovery on data source Name of the datasource has started.	1496
BTAVM0007I Discovery on data source Name of the datasource completed successfully.	1496
BTAVM0008I Probe of hypervisor Name of the Hypervisor has started.	1497
BTAVM0009I Probe of hypervisor Name of the Hypervisor completed successfully.	1497
BTAVM0010I A connection test to data source Name of the data source has started.	1497
BTAVM0011I The Connection test to data source Name of the data source completed successfully.	1497
BTAVM0012I Hypervisor Name of the Hypervisor discovered/rediscovered.	1497
BTAVM0013I Discovery: Hypervisor Name of the hypervisor will not be discovered as it is managed by another data source.	1498
BTAVM0014I Discovery: Hypervisor Name of the hypervisor will not be discovered as it itself is registered as a data source.	1498
BTAVM0015I Collection of the physical storage configuration for hypervisor Name of the hypervisor has started.	1498
BTAVM0016I Collection of the physical storage configuration for hypervisor Name of the hypervisor completed successfully.	1498
BTAVM0017I Collection of the logical storage configuration for hypervisor Name of the hypervisor has started.	1498
BTAVM0018I Collection of the logical storage configuration for hypervisor Name of the hypervisor completed successfully.	1499
BTAVM0019I Collection of the virtual machines configuration for hypervisor Name of the hypervisor has started.	1499
BTAVM0020I Collection of the virtual machines configuration for hypervisor Name of the hypervisor completed successfully.	1499
BTAVM0021I The probe of name of the hypervisor found number of physical disks physical disks.	1499
BTAVM0022I The probe of name of the hypervisor found number of logical volumes logical volumes.	1499
BTAVM0023I The probe of name of the hypervisor found number of virtual machines virtual machines.	1500
BTAVM0024I The Name of the hypervisor hypervisor was removed.	1500
BTAVM0025I VMWare Cluster Name of the Cluster discovered/rediscovered.	1500
BTAVM1301I The probe of name of the hypervisor could collect partial information only for the disk with the device name Device name of the disk.	1500
BTAVM1302I LUN correlation is not supported for disk with device name Device name of the disk, vendor: Vendor name, model: model name, for hypervisor hypervisor name.	1500
BTAVM1503E An internal error occurred: Text describing the internal error.	1501
BTAVM2001E The mandatory parameter Name of the mandatory parameter which is missing is missing.	1501
BTAVM2002E Invalid parameter Name of the parameter which was invalid.	1501
BTAVM2003E A database error was encountered during database query or insert.	1501
BTAVM2004E Cannot connect to the database repository.	1501
BTAVM2006E The operation Name of the operation that failed failed for the following reason: Reason of the failure.	1502
BTAVM2007E The Web service call Name of the operation failed for the following reason: Reason of the failure.	1502
BTAVM2008E The product Name of the unsupported product is not supported.	1502
BTAVM2010E The user name or password is invalid for Address of the host	1502
BTAVM2011E The operation Name of the timed out operation could not complete within the time limit of Timeout threshold in milliseconds milliseconds.	1503
BTAVM2012E An error occurred while trying to establish secure communication over SSL.	1503
BTAVM2013E The Add Device wizard could not add the Name of the data source data source.	1503
BTAVM2014E The deletion of data source Name of the data source failed.	1503
BTAVM2015E The modification of data source Name of the data source failed.	1504
BTAVM2016E Discovery on data source Name of the datasource failed.	1504
BTAVM2017E Probe of the hypervisor Name of the Hypervisor failed.	1504
BTAVM2018E IBM Spectrum Control can't connect to the data source Name of the datasource.	1504
BTAVM2201E Probe: An error occurred during the collection of the physical storage configuration.	1505
BTAVM2202E Probe: An error occurred during the collection of the logical storage configuration.	1505
BTAVM2204E Probe: An error occurred during the collection of the virtual machine configuration.	1505
BTAVM2206E Discovery: the hypervisor Name of the hypervisor will not be discovered because its version is not supported.	1505
BTAVM2207E Calculation of the summary data for the hypervisor Name of the hypervisor failed.	1506
BTAVM2208E Unable to obtain the hypervisor version(s) from the datasource Name of the datasource.	1506
BTAVM2209E Unable to obtain information about other Virtual Centers managing the hypervisor(s) of datasource Name of the datasource.	1506
BTAVM2210W Error getting LUN definition data for the disk with the device name Device name of the disk, storage subsystem vendor: Vendor name, model: model name, for hypervisor hypervisor name.	1506

BTAVM2211E Probe: Virtualization Manager failed to get the VMWare VI data source for the hypervisor Name of the hypervisor from the database.	1507
BTAVM2212E Probe: The hypervisor Name of the hypervisor is not available on the VMWare VI datasource Name of the datasource.	1507
BTAVM2213E Data source Name of the datasource is disconnected from Virtual Center.	1507
BTAVM2214E The probe job encountered an NFS file system while probing ESX server {0}. IBM Spectrum Control currently does not support probes of ESX servers with NFS file systems. The probe job for this ESX server has been stopped. Probes of other ESX servers that are included in this probe job will continue.	1507
BTAVM2215W Unsupported storage subsystem disk with device name Device name of the disk, vendor: Vendor name, model: model name, for hypervisor hypervisor name with hypervisor version less than 3.5.0.	1507
BTAVM2216E Unable to get keystore instance.	1508
BTAVM2217E Unable to load keystore file.	1508
BTAVM2218E Unable to set certificate entry in keystore file.	1508
BTAVM2219E Unable to open keystore for writing.	1508
BTAVM2220E Unable to close keystore file.	1509
BTAVM2221E Unable to acquire lock on keystore file.	1509
BTAVM2222E Unable to store certificate in keystore file.	1509
BTAVM2223E Unable to release lock on keystore file.	1509
BTAVM2224E Unable to decrypt keystore password.	1509
BTAVM2225E Unable to open keystore for reading.	1509
BTAVM2226E Certificate already exists in keystore.	1510
BTAVM2227E host_address hypervisor is already being monitored and could not be added.	1510
BTAVM2228E Missing host name.	1510
BTAVM2229E Missing certificate.	1510
BTAVM2230E Cannot create keystore directory.	1510
BTAVM2231E Cannot download the certificate from Data Source Name of the data source.	1511
BTAVM2232E Cannot connect to the Name of the data source data source.	1511
BTAVM2233E Cannot download the certificate from the port.	1511
BTAVM2234E The hypervisor name hypervisor was not removed because IBM Spectrum Control is running other actions on the device.	1511
BTAVM2235E Unable to obtain the cluster(s) from the datasource Name of the datasource.	1511
BTAVM2236W Subsequent steps of probe process may not be able to collect data for the hypervisor Name of the hypervisor because the hypervisor is in critical state.	1512
BTAVM2237E Datastore Browser Task failed for hypervisor Name of the hypervisor, datastore Name of the datastore with error: Error	1512
BTAVM2238E The registration of the vSphere Web Client extension for IBM Spectrum Control has started on Name of the vCenter server.	1512
BTAVM2239E The registration of the vSphere Web Client extension for IBM Spectrum Control did not extract the extension package.	1512
BTAVM2240E The registration of the vSphere Web Client extension for IBM Spectrum Control did not complete while updating the VASA web archive file, vasa.war, with the IBM Spectrum Control server configuration.	1513
BTAVM2241E The registration of the vSphere Web Client extension for IBM Spectrum Control completed.	1513
BTAVM2242E Unable to register IBM Spectrum Control as an extension on the vCenter server Name of the vCenter server. The validation of input values did not complete.	1513
BTAVM2243E Unable to register IBM Spectrum Control as an extension on the vCenter server Name of the vCenter server. Could not authenticate with the vCenter server.	1513
BTAVM2244E The registration of the vSphere Web Client extension for IBM Spectrum Control did not complete.	1514
BTAVM2245E Unable to connect to the vCenter Server Name of the datasource.	1514
BTAVM2246E Unable to configure the vCenter Server.	1514
BTAVM2247E The registration of the vSphere Web Client extension for IBM Spectrum Control did not delete the temporary directory Name of the directory.	1514
BTAVM2248E The registration of IBM Spectrum Control as a VASA provider did not complete.	1515
BTAVM2249E Automatic registration of IBM Spectrum Control as a VASA provider is not supported for vCenter Server version 5.0 and earlier.	1515
BTAVM2250E IBM Spectrum Control is already registered as a VASA provider for vCenter Server server_name. Register IBM Spectrum Control as a VASA provider manually in the vSphere Web Client to update the credentials.	1515
BTAVM2251E One or more third-party VASA providers are already registered with the vCenter Server. IBM Spectrum Control VASA provider was not registered. Register IBM Spectrum Control as a VASA provider manually.	1515
BTAVM2252E The registration of IBM Spectrum Control as a VASA provider has started on Name of the vCenter server.	1516
BTAVM2253E The registration of IBM Spectrum Control as a VASA provider has completed.	1516
BTAVM2254E The registration of the vSphere Web Client extension for IBM Spectrum Control did not complete. The current session is invalid.	1516
BTAVM2255E The registration of IBM Spectrum Control as a VASA provider did not complete. The current session is invalid.	1516
BTAVM2256W Could not determine the host for VM with ID: host id and Name: Vendor name. Check if the same mac address is used on other computers.	1516
BTAVM2257I Found number of files files on name of datastore of name of the hypervisor.	1517
BTAVM2258I The probe of name of the hypervisor found number of controllers controllers.	1517
BTAVM2259I Collecting file system details for hypervisor Name of the hypervisor.	1517
BTAVM2260I Collecting list of files for hypervisor Name of the hypervisor.	1517
BTAVM2261I Collecting logical volumes for hypervisor Name of the hypervisor.	1517
BTAVM2262I Collecting disk partition for hypervisor Name of the hypervisor.	1518
BTAVM2263I Files details for Name of the datastore being collected by id of the Hypervisor.	1518
BTAVM2264I Files details for Name of the datastore were collected by id of the Hypervisor on timestamp.	1518
BTAVM2265E Invalid host name or IP address.	1518
BTAVM2266E The connection information cannot be updated because it points to another device.	1518
BTAVM2268E The connection information cannot be updated because IBM Spectrum Control cannot determine if the hypervisor is managed by the Name of the data source data source.	1519
BTAVM2269E The connection information cannot be updated because a data source with this host name or IP address is already present.	1519
BTAVM2270E The connection information cannot be updated because it doesn't point to a data source of the same type (vCenter/ESX).	1519
BTAVM2271W The hypervisor Name of the Hypervisor cannot be discovered because its connection state is "Connection State".	1519
BTAVM2272E The user User Name does not have the privilege to browse the datastore Name of the Datastore.	1519
HWN020001I Operation Name of the operation processed successfully.	1520
HWN020002E Mandatory parameter Name of the mandatory parameter which is missing missing	1520
HWN020003E Invalid parameter Name of the parameter which was invalid	1520

HWN020101E The external process terminated unexpectedly.	1520
HWN020102W The external process was canceled per users request.	1520
HWN020103E The external process exceeded the timeout limit and was canceled.	1521
HWN020104E The external process could not be started.	1521
HWN020105E The data collector is not responding to the server.	1521
HWN020106E An external process was cancelled by the data collector.	1521
HWN021503E The action cannot be completed	1521
HWN021504E Entity The ID of the entity was not found.	1522
HWN021508E Credentials not found	1522
HWN021514E The invocation of CIM method Name of method failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method	1522
HWN021515E The invocation of CIM method Name of method failed on SMI-S provider Name of SMI-S provider with the following exception text: Exception text	1522
HWN021516E The LSS specified LSS name on subsystem Name of subsystem is already at the maximum volume number (255). Volume creation can not be done on this LSS, please select a different one.	1523
HWN021517E The connection to SMI-S provider for storage system VPD of the storage system could not be made.	1523
HWN021520E The attribute Name of the attribute was not found.	1523
HWN021522E Host port The WWPN of the host port not assigned to Volume The PK of the volume	1523
HWN021524E Indexed Properties Names don't match	1523
HWN021529E An SMI-S provider has reported unexpected values: IP and port of SMI-S provider.	1524
HWN021530E The Volume - Port mapping can not be created. There are existing mappings that prevent this combination. VolumeCOP: The ID of the volume , Port: The WWPN of the port that should be mapped to the volume	1524
HWN021531E SMI-S provider The IP and port of the SMI-S provider can not reach storage system The VPD of the storage system	1524
HWN021535E There is not enough space left in the storage pool The primary key of the Pool on storage system The VPD of the storage system to create a volume of The requested volume size bytes.	1524
HWN021536E The CIM method The CIM method that is not supported. is not supported on the storage system The VPD of the storage system	1525
HWN021537E Could not create connection to SMI-S provider The IP and port for the SMI-S provider..Reason: The exception returned by the SMI-S provider.	1525
HWN021538E The username The username that was used to connect to the SMI-S provider. or password is wrong on SMI-S provider The IP and port for the SMI-S provider.	1525
HWN021539E The SVC with IP The IP of the SVC. which is managed by SMI-S provider The IP and port for the SMI-S provider. can not be discovered. The status is The status of the SVC. .	1525
HWN021540E The invocation of CIM method Name of method failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method. Details provided by the SMI-S provider : Description of Returncode	1525
HWN021600W Operation Name of the operation. partially processed.	1526
HWN021601E The operation(s) Operation_names failed.	1526
HWN021602E It is necessary to specify target ports for storage device VPD of the storage subsystem	1526
HWN021603W More storage volumes and ports than specified will loose access	1526
HWN021604E WWPNs and storage volumes to be unassigned not completely specified. Assigned WWPNs: All WWPNs that are assigned to the volumes in the host port collection , missing WWPNs: The WWPNs that are assigned but were not specified in the input parameter in the method unassign . Storage volumes to be unassigned not completely specified. Assigned storage volumes: Lists all storage volumes that are really assigned to the WWPNs. }, missing storage volumes: The storage volumes that are really assigned but were not specified in the input parameter in the method unassign	1526
HWN021605I More storage volumes and ports than specified will gain access.	1527
HWN021606E WWPNs and storage volumes to be assigned not completely specified. Missing WWPNs: The WWPNs that need to be assigned but were not specified in the input parameter. . Storage volumes to be assigned not completely specified. Missing storage volumes: The storage volumes that need to be assigned but were not specified in the input parameter.	1527
HWN021607E The client type the client type with description the client description is not supported on SMI-S provider the SMI-S provider IP and port for storage subsystem the subsystem ID of volumes the volumeIDs of the subsystem which were passed in	1527
HWN021608E The target port the target port ID does not belong to storage subsystem the subsystem ID of volumes the volumeIDs of the subsystem which were passed in	1528
HWN021609E There is not enough space left in the storage pool The primary key of the Pool on storage system The VPD of the storage system to create The number of volumes to create volumes of The total size needed bytes total.	1528
HWN021610E The specified size The size of the volume to create is not supported on pool The storage pool ID Size has to be dividable by Divisor returned by getSupportedSizeRange and in between Minimum returned by getSupportedSizeRange and Maximum returned by getSupportedSizeRange	1528
HWN021611E Volume The volume ID has mappings, it can not be deleted.	1528
HWN021612E The mapping between volume The volume ID and port The initiator port wwpn exists already	1528
HWN021613E The WWPN The WWPN not found can not be found on subsystem The subsystem	1529
HWN021614E The WWPNs The WWPNs without mappings have no mappings on storage system The storage system	1529
HWN021615E WWPNs WWPNs that can not share mappings can not share mappings on storage system Storage system}. There are existing mappings that prevent this.	1529
HWN021616E Volumes VolumeIDs can not share mappings on storage system Storage system }. There are existing mappings that prevent this.	1529
HWN021617E The stored data for storage system The storage system is not in sync with the environment. Rerun data collection.	1530
HWN021618E Modifying target ports is not supported by subsystem the subsystem .	1530
HWN021619E Modifying the target ports for mapping of initiator port initiator port WWPN and volume volume name will also modify the target ports of the following mappings: port - volume list	1530
HWN021620I Modifying the target ports for mapping of initiator port initiator port WWPN and volume volume name will modify the target ports of more mappings than specified.	1530
HWN021621E It is not supported to modify the target ports of existing mappings and create new mappings in one step. Modify the existing mappings first and then create the new mappings. Existing mappings: port - volume list	1530
HWN021622I Started modification of the assignment of volume VolumeID on subsystem Subsystem to initiator port WWPN . Target ports to add: target ports to add Target ports to remove: target ports to remove	1531
HWN021623I Finished modification of the assignment of volume VolumeID on subsystem Subsystem to initiator port WWPN . Target ports to add: target ports added Target ports to remove: target ports removed	1531
HWN021624E The modification of the assignment of volume VolumeID on subsystem Subsystem to initiator port WWPN failed. Target ports to add: target ports to add Target ports to remove: target ports to remove	1531
HWN021650E A timeout occurred while connecting to SMI-S provider SMI-S provider IP and port.	1531

HWN021651E Job on SMI-S provider SMI-S provider IP and Port in format IP:Port failed. Job Status: Job status . Error code is Error code , error description: Error description . Check IBM Spectrum Control and SMI-S provider logs.	1532
HWN021652E The process has timed out. Check the IBM Spectrum Control log files for more information.	1532
HWN021653E The attribute Name of the attribute was not found.	1532
HWN021654E Pool ID was not found.	1532
HWN021655E Volume ID The ID of the volume was not found.	1532
HWN021656E Port ID The ID of the port was not found.	1533
HWN021657E Subsystem ID The ID of the subsystem was not found.	1533
HWN021658E Managed Disk ID The ID of the MDisk was not found.	1533
HWN021659E SMI-S provider The ID of the SMI-S provider was not found	1533
HWN021660E IO Group The SVC IO Group was not found.	1533
HWN021661E Extent The storage extent external key was not found.	1533
HWN021662E Physical volume The physical volume external key was not found.	1534
HWN021670E The client type the client type with description the client description is not unique on SMI-S provider the SMI-S provider IP and port } for storage subsystem the subsystem ID of volumes the volumeIDs of the subsystem which were passed in	1534
HWN021671I The storage system The storage system was deleted from the database	1534
HWN021672E The storage system name storage system was not removed because other monitoring actions are running on the device.	1534
HWN021673E The probe job on SMI-S provider SMI-S provider IP and Port in format IP:Port did not complete within the time limit of Microseconds microseconds. The job is Percent complete percent complete. Check the SMI-S provider log for job status. Job information: JobCOP . Run the probe job again after the current job has completed.	1534
HWN021674E Job on SMI-S provider SMI-S provider IP and Port in format IP:Port returned unexpected results. Job information: JobCOP Job status: JobState , status description: JobStatus Check SMI-S provider log. Redo probe if the job completed.	1535
HWN021675I Started creation of volume with size Size in pool Pool on subsystem Subsystem	1535
HWN021676I Volume creation completed successfully. New volume VolumeID created with size Size in pool Pool on subsystem Subsystem .	1535
HWN021677E Volume creation failed. The volume of size Size in pool Pool on subsystem Subsystem could not be created.	1535
HWN021678I Started assignment of volume VolumeID on subsystem Subsystem to initiator port WWPN .	1536
HWN021679I Finished assignment of volume VolumeID on subsystem Subsystem to initiator port WWPN .	1536
HWN021680E The assignment of volume VolumeID on subsystem Subsystem to initiator port WWPN failed.	1536
HWN021681I Started unassignment of volume VolumeID on subsystem Subsystem to initiator port WWPN .	1536
HWN021682I Finished unassignment of volume VolumeID on subsystem Subsystem to initiator port WWPN .	1536
HWN021683E The unassignment of volume VolumeID on subsystem Subsystem to initiator port WWPN failed.	1537
HWN021684I Started deletion of volume VolumeID on subsystem Subsystem .	1537
HWN021685I Volume deletion completed successfully. Volume VolumeID on subsystem Subsystem was deleted.	1537
HWN021686E Volume deletion failed. Volume VolumeID on subsystem Subsystem could not be deleted.	1537
HWN021687I Started modification of Pool Pool display name on subsystem Subsystem display name .	1537
HWN021688I Pool modification completed successfully. Pool Pool display name on subsystem Subsystem display name was modified.	1538
HWN021689E Pool modification failed. Pool Pool display name on subsystem Subsystem display name could not be modified.	1538
HWN021690I Started creation of number volumes volumes with size Size in pool Pool on subsystem Subsystem	1538
HWN021691I Created number volumes out of total number volumes volumes with size Size in pool Pool on subsystem Subsystem	1538
HWN021692E Volume creation failed. Created number volumes out of total number volumes volumes with size Size in pool Pool on subsystem Subsystem	1538
HWN021693W Warning: The task succeeded, but the database update failed. Run probe to update the database.	1539
HWN021700I Enumerating CIM Associator The CIM association name which is being enumerated. for The name of the DB table which will be populated as result of this query.	1539
HWN021701I Enumerating CIM Class The CIM class name which is being enumerated. for The name of the DB table which will be populated as result of this query.	1539
HWN021702I Querying SMI-S provider	1539
HWN021703I Task starting on SMI-S provider Identifier of the SMI-S provider..	1539
HWN021708I Initializing Collection for storage system storage system identification.	1540
HWN021709I Collection for storage system storage system identification completed.	1540
HWN021710I Discovering devices for SAN Volume Controller The VPD of the SAN Volume Controller.	1540
HWN021711I Discovery devices for SAN Volume Controller The VPD of the SAN Volume Controller. failed with error message The exception which has occurred	1540
HWN021712I Collecting Nodes for storage system storage system identification.	1540
HWN021713I Collecting fibre channel ports for storage system storage system identification.	1541
HWN021714I Collecting volumes for storage system storage system identification.	1541
HWN021715I Traversing host to volume assignments for storage system storage system identification.	1541
HWN021716I Collecting pools and volumes for storage system storage system identification.	1541
HWN021717I Collecting volume settings for storage system storage system identification.	1541
HWN021718I Collecting client setting data for storage system storage system identification.	1542
HWN021719I Perform collection post process tasks for storage system storage system identification.	1542
HWN021720I Flash enclosure is missing drive flash_drive_ identifier.	1542
HWN021724W SMI-S provider SMI-S provider identifier manages device(s) of type device_type which is supported through the native device interface or SNMP only.	1542
HWN021725I IBM Spectrum Control discovered/rediscovered a device with name Identifier of the device. on SMI-S provider Identifier of the SMI-S provider..	1542
HWN021726I IBM Spectrum Control discovered/rediscovered no device on SMI-S provider Identifier of the SMI-S provider..	1543
HWN021727I IBM Spectrum Control discovery starting on SMI-S provider Identifier of the SMI-S provider..	1543
HWN021728I IBM Spectrum Control discovery on SMI-S provider Identifier of the SMI-S provider. is complete.	1543
HWN021729W IBM Spectrum Control discovery of Device type value is not supported.	1543
HWN021730W IBM Spectrum Control discovery of device value with code level value is not supported on SMI-S provider Identifier of the SMI-S provider..	1543
HWN021731I Probing Volumes for Storage System: value.	1544
HWN021732I Number of Volumes Found Currently: value. Continuing to Probe Volumes.	1544
HWN021733I value Volumes Found.	1544
HWN021734I Probing Disks for Storage System: value.	1544

HWN021735I Number of Disks Found Currently: value. Continuing to Probe Disks.	1544
HWN021736I value Disks Found.	1545
HWN021737I Probing Virtual Disks for Cluster: value	1545
HWN021738I Number of Virtual Disks currently found: value. Continuing to probe Virtual Disks.	1545
HWN021739I value Virtual Disks found.	1545
HWN021740I Probing Views of Host Initiator access to Volumes.	1545
HWN021741I value Views Found.	1545
HWN021742E The SMI-S provider SMI-S provider URL is not managing storage subsystems.	1546
HWN021743E The SMI-S provider SMI-S provider URL is not managing switches.	1546
HWN021744E Cannot connect to a resource because of an SSL certificate error. Troubleshooting information: http://www.ibm.com/support/docview.wss?uid=swg21976237	1546
HWN021745I Cannot connect to a resource because of an SSL certificate error. Troubleshooting information: http://www.ibm.com/support/docview.wss?uid=swg21976237. An alternate resource will be used.	1546
HWN021746W SMI-S provider Identifier of the SMI-S provider. manages Cisco device types through SNMP only.	1547
HWN021747E Unable to add the specified switch by using SNMP. The switch is a Brocade switch and can be added only by using an SMI agent.	1547
HWN021800E Failed to get a database connection.	1547
HWN021801E The server failed to get SMI-S provider entity from database.	1547
HWN021802E Experienced SQL problems while working with database: The SQL error.	1547
HWN021803W The server did not get userid and or password for SMI-S provider The Service URL of the SMI-S provider from database.	1548
HWN021804E The server failed to access slp attributes for SMI-S provider The Service URL of the SMI-S provider from database.	1548
HWN021805E CIMOMManager failed to get a database mapper of type The type of the database mapper.	1548
HWN021806E CIMOMManager failed to get a valid mapper result from The type of the database mapper.	1548
HWN021807E CIMOMManager failed to get a proxy for calling slp discovery.	1548
HWN021808E The device cannot be contacted through any of the following SMI-S providers The comma separated list of IP and port for the SMI-S providers.. Possible causes are that the SMI-S providers are not accessible or the device is disconnected from the SMI-S providers.	1549
HWN021809E The host for SMI-S provider The service URL of the SMI-S providers. was not resolvable in DNS.	1549
HWN021810E The service URL for SMI-S provider The service URL of the SMI-S providers. is not valid.	1549
HWN021811I The operational status for device The ID of the device. on SMI-S provider The service URL of the SMI-S provider. has this value The operational status vector. .	1549
HWN021812E The operational status for device The ID of the device. on SMI-S provider The service URL of the SMI-S provider. could not be retrieved because SMI-S provider is in status The SMI-S provider connection status. .	1549
HWN021813E Fabric ID The ID of the fabric was not found.	1550
HWN021814E The device device id cannot be contacted through the SMI-S provider SMI-S provider service URL.	1550
HWN021899E Switch The wwn of the switch. has no associated Fabric.	1550
HWN021901E The virtual disk size cannot exceed maximum size when creating space efficient virtual disks.	1550
HWN021902E Invalid grain size. Valid values are valid values.	1550
HWN021903E Authentication to ip or name of host failed. Please specify correct authentication information.	1551
HWN021904E Connection to IP address or name of host failed with following operating system exception: exception text . Please make sure IP address is correct and machine is up and running. If this is a SVC V4 machine, it could be that its RAS interface is not up. If this is a SVC V5, make sure the SMI-S provider is up and running.	1551
HWN021905E Connection to IP address or name of host failed with following operating system exception: exception text .	1551
HWN021906E Failed to get native API entity from database.	1551
HWN021907E The IP address The service URL of the SMI-S providers. was not resolvable in DNS.	1552
HWN021908E Failed to get a proxy for calling NAPI discovery.	1552
HWN021909E There are no IO Groups available for Virtual Disk creation.	1552
HWN021910E Managed Disk ID The ID of the MDisk is not in unmanaged mode and cannot be added to the specified managed-disk group.	1552
HWN021911E Another probe of storage subsystem The Name+Nameformat of the storage subsystem is already in progress.	1552
HWN021912E Other probes of storage subsystems The list of Name+Nameformat of the storage subsystems are already in progress.	1553
HWN021913E IBM Spectrum Control Device Server could not write to directory The directory.	1553
HWN021914E SSH key file The SSH key file name is still in use, so it cannot be deleted.	1553
HWN021915E IBM Spectrum Control Device Server could not delete the file The file.	1553
HWN021916E The storage subsystem subsystem ID is not configured for file level management.	1553
HWN021917E An invalid parameter Name of the parameter which was invalid was specified. The corresponding file system mount point does not exist.	1554
HWN021919E The cluster ID The ID of the cluster. was not found.	1554
HWN021920E The export ID The ID of the export. was not found.	1554
HWN021921E The specified activity or protocol could not be used to change the export The ID of the export..	1554
HWN021922E The file system ID file_system_ID was not found.	1554
HWN021923E Invalid parameter Name of the parameter which was invalid. File system does not exist.	1554
HWN021924E The parameter Name of the parameter which was invalid is not a valid parameter.	1555
HWN021925E The fileset ID fileset_ID was not found.	1555
HWN021926E The WAN-cache source ID WAN_cache_source_id was not found.	1555
HWN021927E The WAN-cache ID WAN_cache_source_id was not found.	1555
HWN023000I The Optimization Execution task has started.	1555
HWN023001E The task to optimize the volumes was not completed successfully.	1556
HWN023002I The Optimization Execution task has completed.	1556
HWN023003I The Optimization Execution task retrieved number recommendations	1556
HWN023004I The Optimization Automation request persisted recommendations to be processed.	1556
HWN023005I The Optimization Execution task updated the status of number recommendations.	1557
HWN023006I The Optimization Automation request begins processing number recommendations.	1557
HWN023007W The recommendation being processed contains a virtual disk that is no longer detected.	1557
HWN023008W The recommendation for virtual disk vdisk name contains a source storage pool that is no longer detected.	1557
HWN023009W The recommendation for virtual disk vdisk name contains a target storage pool that is no longer detected.	1557

HWN023010I Virtual disk vdisk name was successfully migrated from storage pool source pool name to storage pool target pool name.	1558
HWN023011W The recommendation for virtual disk vdisk name contains a virtual disk that does not exist in the source storage pool source pool name or the target storage pool target pool name.	1558
HWN023012W The recommendation for virtual disk vdisk name contains a non-mirrored virtual disk that is now a mirrored virtual disk.	1558
HWN023013W The recommendation for virtual disk vdisk name contains a mirrored virtual disk that is now a non-mirrored virtual disk.	1558
HWN023014I The recommendation for virtual disk vdisk name requires more space on target pool target pool name to be processed.	1559
HWN023015I Virtual disk vdisk name will now be migrated from storage pool source pool name to storage pool target pool name.	1559
HWN023016I Successfully added virtual disk copy to virtual disk vdisk name.	1559
HWN023017I Synchronization for virtual disk vdisk name has completed synchronization percent% and requires about seconds to complete seconds to complete.	1559
HWN023018I Synchronization for virtual disk vdisk name has completed.	1559
HWN023019I Successfully removed a virtual disk copy from virtual disk vdisk name.	1560
HWN023020I Successfully changed the synchronization rate of virtual disk vdisk name to syncrate%.	1560
HWN023021I Successfully changed the primary copy of virtual disk vdisk name.	1560
HWN023022E There is no space available on target pool target pool name to migrate the virtual disk vdisk name.	1560
HWN023023E Unable to submit request to add vdisk copy command for virtual disk vdisk name due to rc (rc).	1560
HWN023024E Unable to complete request to add vdisk copy command for virtual disk vdisk name due to rc (rc).	1561
HWN023025E Unable to submit request to get vdisk synchronization progress for virtual disk vdisk name due to rc (rc).	1561
HWN023026E Unable to complete request to get vdisk synchronization progress for virtual disk vdisk name due to rc (rc).	1561
HWN023027E Unable to submit request to remove vdisk copy command for virtual disk vdisk name due to rc (rc).	1561
HWN023028E Unable to complete request to remove vdisk copy command for virtual disk vdisk name due to rc (rc).	1562
HWN023029E Unable to submit request to change the synchronization rate for virtual disk vdisk name due to rc (rc).	1562
HWN023030E Unable to complete request to change the synchronization rate for virtual disk vdisk name due to rc (rc).	1562
HWN023031E Unable to submit request to change the primary copy for virtual disk vdisk name due to rc (rc).	1562
HWN023032E Unable to complete request to change the primary copy for virtual disk vdisk name due to rc (rc).	1563
HWN023033E The request failed. Message from failed request: message.	1563
HWN023034E The Optimization Automation job completed with errors in the recommendations.	1563
HWN023035W The Optimization Execution task completed with warnings.	1563
HWN023036E The request failed because there were not enough extents in the storage pool.	1564
HWN023037E The request failed because the number of copies of this volume would exceed the limit.	1564
HWN023038E The request failed because the copy specified does not exist.	1564
HWN023039E The following exception occurred during a migration request: exception	1564
HWN023040E The migration request for volume vdisk name is already being processed.	1565
HWN023041W The request to migrate the mirrored volume vdisk name is suspended because the secondary volume is offline.	1565
HWN023042E The secondary copy needed for migration does not exist.	1565
HWN023043I The mirrored volume migration for volume vdisk name will be ignored.	1565
HWN023044I The mirrored volume migration for volume vdisk name will result in the current secondary volume becoming the primary volume.	1565
HWN023045I The mirrored volume migration for volume vdisk name will result in the primary volume being migrated to the target pool.	1566
HWN023046I The Migration of the previously abandoned Optimization Automation job has started.	1566
HWN023047I The Migration of the previously abandoned Optimization Automation job has completed.	1566
HWN023048I The Optimization Automation cancellation job jobname has started.	1566
HWN023049E The Optimization Automation cancellation job completed with errors.	1566
HWN023050I The Optimization Automation cancellation job jobname has completed.	1567
HWN023051I The Optimization Automation job jobname will be canceled.	1567
HWN023052W The Optimization Automation job is not in progress.	1567
HWN023053I The migration of volume vdisk name has been canceled.	1567
HWN023054W The Optimization Automation job was canceled.	1567
HWN023055I The volume that was chosen for transformation, vdisk name, is a secondary volume in a mirrored volume relationship. The secondary volume will be migrated to the specified target pool or converted as specified.	1568
HWN024000I An optimization analysis task was started.	1568
HWN024001I The analysis is completed.	1568
HWN024002W Unable to retrieve any policy for Tier value.	1568
HWN024003I Analyzed number_of_volumes volumes on tier tier_number for storage virtualizer subsystem_name.	1568
HWN024006W No target pools in subsystem value were selected.	1568
HWN024011W Destination storage pool value in subsystem value was not considered. Reason: value.	1569
HWN024012I It is recommended that number_of_volumes volumes on tier source_tier_number are moved to tier target_tier_number.	1569
HWN024015I The optimization analysis of the value subsystem was started.	1569
HWN024016W Volume value is already in the destination storage pool value. No recommendations will be generated for the volume.	1569
HWN024018W No destination storage pools in Tier value have been specified for subsystem value.	1569
HWN024019W The following pools on tier tier_number on the storage_system storage system cannot be balanced by redistributing or re-tiering volumes: pool_names.	1570
HWN024020I Started analysis to balance pools on tier value.	1570
HWN024021W The pool_name pool on tier tier_number on the storage_system storage system cannot be balanced by redistributing the volumes.	1570
HWN024027I Storage Pool pool name has insufficient available space for volume volume name in storage pool pool name.	1570
HWN024030W One or more entities specified as input for the analysis could not be found or pools or volumes in some input entities could not be found.	1571
HWN024031W One or more entities specified as candidate destinations for the analysis could not be found.	1571
HWN024032W For one or more mirrored volumes, both the primary and the secondary volume copies were chosen for transformation. You cannot transform both volume copies in the same transform task. Only the primary volume copies are included for transformation. You can transform the secondary volume copies in a separate transformation.	1571
HWN024033W The volume volume name cannot be analyzed because it is not in a capacity pool.	1571
HWN024034W The pool pool name cannot be analyzed because the pool is not in a capacity pool.	1572
HWN024035W The storage virtualizer system name cannot be analyzed because the storage virtualizer is not in a capacity pool.	1572

HWN024036W The operation to transform the volumes on the subsystem name storage virtualizer cannot be completed because the destination pools were not available.	1572
HWN024037E An unexpected error occurred. The operation to transform the volumes on the subsystem name storage virtualizer cannot be completed because the destination pools were not identified.	1572
HWN024043I The capacity pools of the source volumes were selected as the target pools.	1572
HWN024046I The option that was selected to handle volumes with mirrored volumes is: After optimization, set the copy of the secondary volume in the destination pool as the primary volume. The original secondary volume remains the secondary volume.	1573
HWN024047I The number of days for collecting performance data to analyze the volumes is set to performance_data_collection_period.	1573
HWN024050I Automatic tiering was selected to tier the volumes.	1573
HWN024051I The tiering analysis is starting.	1573
HWN024052I Tier tier# has an I/O density threshold value of value per second per GiB.	1573
HWN024053I Tier tier#, has a file age threshold value of value percent of files last accessed within time_unit.	1574
HWN024054I The real capacity for the thin-provisioned volumes is set to value unit.	1574
HWN024055I The auto expand property of the thin-provisioned volumes is set to yes/no.	1574
HWN024056I The warning level for thin-provisioned volumes is set to value %.	1574
HWN024057I The grain size that was specified for the thin-provisioned volumes is grain_size KiB.	1574
HWN024058I The real capacity for the compressed volumes is set to value unit.	1575
HWN024059I The auto expand property for the compressed volumes is set to yes/no.	1575
HWN024060I The warning level for the compressed volumes is set to value.	1575
HWN024061I The option that was selected to handle volumes with mirrored volumes is: After optimization, set the secondary volume as the primary volume. The volume in the destination pool is the secondary volume.	1575
HWN024062I The option that was selected for mirrored volumes is: Do not optimize volumes with mirrored volumes.	1575
HWN024066I Tier tier# has an I/O rate threshold value of value I/O per second.	1576
HWN024067W Recommendations cannot be generated for number_of_volumes volumes because the volumes do not meet the tiering criteria for tier current_tier_number or for any lower tier.	1576
HWN024068W Recommendations cannot be generated to move number_of_volumes volumes from source_tier to tier target_tier_number due to the pool activity limit value.	1576
HWN024069W Recommendations cannot be generated to move number_of_volumes volumes from tier source_tier to tier target_tier_number because the destination storage pools do not have enough space.	1576
HWN024070I The analysis to optimize subsystem storage_subsystem was completed.	1577
HWN024071I The option that was selected was to restrict the placement of volumes in capacity pools to destination storage pools in the same capacity pool.	1577
HWN024072W No file age information for volume volume name.	1577
HWN024073W Storage pool {0} in tier {1} needs at least one additional storage pool in the same tier for the Balance Analysis to run on this tier.	1577
HWN024074W Storage pool {0} in tier {1} and capacity pool {2} needs at least one additional storage pool in the same tier and capacity pool for the Balance Analysis to run within this capacity pool and on this tier.	1577
HWN024075W number_of_volumes volumes from storage pool pool could not be moved to the destination storage pools because the destination storage pools do not have enough space.	1578
HWN024076W number_of_volumes volumes from storage pool pool could not be moved to the destination pools because the destination storage pools are not in the same capacity pool.	1578
HWN024077W number_of_volumes volumes from storage pool pool could not be moved to the destination storage pools because the destination storage pools would have exceeded the pool activity limit value.	1578
HWN024078W number_of_volumes volumes from storage pool pool could not be moved to the destination storage pools because the destination storage pools already have a volume copy.	1578
HWN024079W Because of an internal error, the number of volumes in the pool storage pool that could not be moved to destination storage pools is number_of_volumes.	1579
HWN024080W Destination storage pool pool already contains a copy of storage volume volume.	1579
HWN024081W Because the destination storage pool does not have sufficient available space, the volume storage volume in the source_pool storage pool cannot be moved to the destination_pool destination storage pool.	1579
HWN024082W Because the destination storage pool contains a copy of the mirrored volume, the volume volume in the source_pool storage pool cannot be moved to the destination_pool destination storage pool.	1579
HWN024083W Because of an internal error, the volume storage volume in the spool storage pool could not be moved to the destination_pool destination storage pool.	1579
HWN024084W Because the destination storage pools contain one or more copies of the mirrored volumes, the number of volumes that could not be moved from tier source_tier to tier target_tier is number_of_volumes.	1580
HWN024085W The pool_name storage pool cannot be balanced because the tier level of the pool was reset to none.	1580
HWN024086E Recommendations cannot be generated because the tier level of the destination_pool_name destination storage pool was reset to none.	1580
HWN024087W Recommendations cannot be generated for one or more of the volumes because collocated volumes cannot be placed in the same destination storage pool.	1580
HWN024088I The option to collocate volumes that are assigned to the same server or hypervisor was selected.	1581
HWN024089I The option to collocate volumes that are assigned to the same server or hypervisor was not selected.	1581
HWN024090W Because the storage pools do not meet the service class requirements, the number of volumes that cannot be moved is no_volumes.	1581
HWN024091W If the recommendation to move the volume_name volume to the storage_pool_name storage pool is implemented, the service class requirements of the volume_name volume cannot be met.	1581
HWN024092W Recommendations cannot be generated to move number_of_volumes volumes from tier source_tier to tier target_tier_number because the destination storage pools do not meet the service class requirements of the volumes.	1581
HWN024093I The number of volumes on tier tier_level that were not analyzed because of the instruction to exclude mirrored volumes from the analysis is number_of_volumes volumes.	1582
HWN024094W Valid target pools were not selected for the subsystem name storage virtualizer.	1582
HWN024095I The grain size for the thin-provisioned volumes was set to the default value of grain_size KiB.	1582
HWN024096W Volumes in the pool_name pool on tier tier_level cannot be moved to a higher tier to reduce the activity level of the pool to the user-defined level.	1582
HWN024097W Volumes in the pool_name pool on tier tier_level cannot be moved to a lower tier to reduce the activity level of the pool to the user-defined level.	1582
HWN024098W Cannot generate recommendations to tier volumes from the storage_system_name storage system because all of the source volumes are in the selected destination storage pools.	1583

HWN024099I The number of volumes that were excluded from the analysis to plan the tiering of the storage_system_name storage system is vols_count. The volumes were excluded because performance data is not available for the volumes.	1583
HWN024100I The number of volumes that were excluded from the analysis to plan the tiering of the storage_system_name storage system is vols_count. The volumes were excluded from the analysis because the capacity of the volumes is zero.	1583
HWN024101I The number of volumes that were excluded from the analysis to plan the tiering of the storage_system_name storage system is vols_count. The volumes were excluded from the analysis because the volumes are not assigned to pools that are tiered or the thresholds were not defined for the tiers.	1583
HWN024102W The recommendation to move the storage_volume_name volume from the source_pool_name storage pool to the target_pool_name storage pool was not generated because the status of the destination pool is offline or excluded.	1584
HWN024103I Reclaiming volumes	1584
HWN024104I Planning for tiering volumes	1584
HWN024105W The recommendation to move the storage_volume_name volume from the source_pool_name storage pool to the target_pool_name storage pool will not be executed because the status of the destination pool is offline or excluded.	1584
HWN024106W The recommendation to move the storage_volume_name volume from the source_pool_name storage pool was not generated because the status of the volume is offline.	1584
HWN024107W The recommendation to move the storage_volume_name volume from the source_pool_name storage pool to the target_pool_name storage pool will not be executed because the status of the volume is offline.	1585
HWN024108E The recommendations can't be shown because the analysis was not completed.	1585
HWN024109W The data for the previous analysis of the storage_subsystem storage system was not deleted.	1585
HWN024110E Volumes reclamation analysis failed for storage_subsystem storage subsystem.	1586
HWN024111W Recommendations cannot be generated to move number_of_volumes volumes from tier source_tier to tier target_tier_number because there is no potential destination pool assigned to the recommended tier.	1586
HWN024112W Cannot generate recommendations to tier volumes from the storage_system_name storage system because the source storage pools and the selected destination storage pools are assigned to the same tier.	1586
HWN024200I The days of the week to include in the analysis: days_of_week.	1586
HWN024201I The time window for the performance data to include in the analysis is set to start time - end time.	1586
HWN024202I The time window for the performance data to include in the analysis is set to start time - end time. The end time occurs on the next day.	1587
HWN024203W The volume storage_volume_name cannot be converted or moved because the target pools do not have sufficient available space or the target pool types are incorrect for the operation.	1587
HWN025000I Storage pool value in storage system value has storage from different types of back-end storage systems. Back-end disk data cannot be determined.	1587
HWN025001I Storage pool value in storage system value has storage from unknown back-end storage system(s). Back-end disk data cannot be determined.	1587
HWN025002I Storage pool value in storage system value has storage from multiple back-end storage systems or from multiple pools in a single storage system. Back-end disk data cannot be determined.	1588
HWN025003I Storage pool value in storage system value has storage from a back-end storage pool with multiple disk types. Back-end disk data cannot be determined.	1588
HWN025004I Storage pool value in storage system value has storage from a back-end storage pool with a mixed raid type. Back-end disk data cannot be determined.	1588
HWN025005I Storage pool value in storage system value has storage from a back-end storage pool with multiple raid types. Back-end disk data cannot be determined.	1588
HWN025006I Storage pool value in storage system value has storage from back-end disks of unknown type. Back-end disk data cannot be determined.	1588
HWN025007I Storage pool value in storage system value has storage from unknown number of back-end disks. Back-end disk data cannot be determined.	1589
HWN025008I Storage pool value in storage system value has storage from back-end disks with unknown raid type. Back-end disk data cannot be determined.	1589
HWN025009E Connection to Data Server failed. Make sure Data Server is up.	1589
HWN025011W All of the target ports for the storage system are used for the provisioning request. The request might take a long amount of time.	1589
HWN025010I Collecting parent pool volumes for storage system: storage system identification.	1590
HWN025011E The port the target port ID has a usage restriction which prevents it from being used as a target port for volume assignment.	1590
HWN025012E The invocation of CIM method ExposePaths failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method.	1590
HWN025013E The invocation of CIM method HidePaths failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method.	1591
HWN025014E The invocation of CIM method CreateOrModifyElementFromStoragePool failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method.	1591
HWN025015E The invocation of CIM method ReturnToStoragePool failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method.	1592
HWN025016E The invocation of CIM method DeleteStorageHardwareID failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method.	1592
HWN025017E A CLI command failed. Check the logs from EP working dir.	1592
HWN025018E An error occurred when attempting to parse the file File name.	1593
HWN025019E The requested operation failed. Check the logs from EP working dir.	1593
HWN025020E The volume cannot be created. The volume of size Size in pool Pool on storage system Subsystem cannot be created. The pool might already have the maximum number of volumes allowed.	1593
HWN025021E Unable to resolve the address for the device because the request was not processed by the data collector.	1593
HWN025022E The data collection detected storage system New Subsystem with serial number new serial number instead of expected serial number expected serial number.	1593
HWN025025I Starting the task to send the report for schedule Schedule Id by email.	1594
HWN025026I The report title report is being created.	1594
HWN025027I The report title report with ID report id is being sent by email to the reports recipients.	1594
HWN025028I The report title report with ID report id was sent by email to the reports recipients.	1594
HWN025029E Can't retrieve the configured settings of the report for schedule Schedule Id .	1594
HWN025030E The report can't be sent because the email server was not configured.	1594
HWN025031E Can't send the report title report with ID report id by email because of the following error: reported_error.	1595
HWN025031I To view the report, choose HTML as the message format or use an email application that supports HTML message formats.	1595
HWN025032E Job failed during post processing of collected data from the data source.	1595
HWN025033E Failed to send the report name report for schedule Schedule Id.	1595
HWN025034I Created number_of_servers agentless servers automatically.	1595
HWN025035I Removed number_of_servers agentless servers automatically.	1596
HWN025036E Can't save the report in the directory.	1596

HWN025037E Can't save the report because the path specifies a file name instead of a directory name.	1596
HWN025038E Can't save the report, because the directory doesn't exist.	1596
HWN025039E Can't save the report because the directory doesn't have enough disk space.	1596
HWN025040I The report title report with ID report id is being saved as report file name in the full path directory.	1597
HWN025041I The report title report with ID report id was saved as report file name in the full path directory.	1597
HWN099990I The method name of the Device Server method of the device server returned return value @(execution context information).	1597
HWN099991I info trace message@(execution context information)	1597
HWN099992W warning trace message@(execution context information)	1597
HWN099993E error/exception trace message @(execution context information)	1597
HWN099994I An object of class name of the class has been instantiated @(execution context information).	1598
HWN099995I === class name.method name entry, parameter(s): parameter value(s) @(execution context information).	1598
HWN099996I === class name.method name exit, return value: method return value (execution time in milliseconds) @(execution context information).	1598
HWN099997I External service name of the (DM) external service will be invoked with parameter(s) parameter value(s)@(execution context information).	1598
HWN099998I Invocation of external service name of the (DM) external service returned result invocation result@(execution context information).	1598
HWN099999I The method name of the device server method of the device server was invoked with parameters invocation parameters@(execution context information).	1599
HWN200000I Probe of switch switch_name completed successfully.	1599
HWN200001I Started post-processing tasks after data was collected for switch switch_name.	1599
HWN6001I Operation operation completed successfully.	1599
HWN6002I Unable to set up NLS message file processing.	1599
HWN6003E Unable to set up tracing.	1600
HWN6004E Operation operation failed.	1600
HWN6005E Unknown operation operation.	1600
HWN6006E Could not initialize connection, rc is rc	1600
HWN6007E Could not parse command arguments: arg	1600
HWN6008E Error processing command: command	1600
HWN6009E Missing 'operation' property in input file	1601
HWN6010I Task arg completed successfully	1601
HWN6011E Task arg failed	1601
HWN6012E Cannot connect to this IP, switching to IP	1601
HWN6013E An IBM XIV CLI command failed. The error is arg.	1601
HWN6014I Command arg completed successfully	1602
HWN6015E Command command failed.	1602
HWN6016I Connected with IP address IP	1602
HWN6017I Started creation of volume with size size in pool pool.	1602
HWN6018I Volume creation completed successfully. New volume volume created with size size in pool pool.	1602
HWN6019I Started deletion of volume volume in pool pool.	1602
HWN6020I Volume deletion completed successfully. Volume wolume deleted in pool pool	1603
HWN6021I Started creation of host host with initiator ports ports	1603
HWN6022I Finished creation of host host with initiator ports ports	1603
HWN6023I Started assignment of volume volume to host host.	1603
HWN6024I Finished assignment of volume volume to host host.	1603
HWN6025I Started unassignment of volume volume from host host.	1603
HWN6026I Finished unassignment of volume volume from host host	1604
HWNEP0001I Successfully persisted number of count instances.	1604
HWNEP0002E The probe failed as the data collector couldn't write to its output file, value.	1604
HWNEP0003E A DS8000 ESSNI command failed. The error code is error_code.	1604
HWNEP0004I Started creation of volume group volume_group.	1605
HWNEP0005I Finished creation of volume group volume_group with subsystem volume group number number .	1605
HWNEP0006I Started adding volumes, with serial numbers volume_list, to subsystem volume group volume_group_number .	1605
HWNEP0007I Finished adding volumes to volume group.	1605
HWNEP0008I Started assignment of host host on subsystem subsystem to volume group volume_group.	1605
HWNEP0009I Finished assigning host on subsystem subsystem to volume group volume_group.	1605
HWNEP0010I Started removing volumes, with serial numbers volume_list, from subsystem volume group volume_group_number .	1606
HWNEP0011I Finished removing volumes, with serial numbers volume_list, from subsystem volume group volume_group_number .	1606
HWNEP0012I Increased virtual capacity of storage pool storage_pool on subsystem subsystem to size size .	1606
HWNEP0013I Collecting pools for storage system storage system identification.	1606
HWNEP0014I Collecting volumes for lss logical subsystems on storage system storage system identification.	1606
HWNEP0015I Collecting volume groups on storage system storage system identification.	1607
HWNEP0016I Collecting hosts on storage system storage system identification.	1607
HWNEP0017I value Hosts Found.	1607
HWNEP0018I Launching external process for devices devices.	1607
HWNEP0019I External process for devices devices completed successfully.	1607
HWNEP0020E Could not create connection to NAPI The IP for the NAPI..	1608
HWNEP0021E ESSNI API query for Space Efficient Volume failed with ESSNI code ESSNI Code. Data from ESSNI is considered suspect.	1608
HWNEP0022I Started deletion of volume group with number volume_group_number.	1608
HWNEP0023I Finished deletion of volume group with number volume_group_number.	1608
HWNEP0100I Probing Volumes for Storage System: value	1608
HWNEP0101I Number of Volumes currently found: value. Continuing to probe Volumes.	1609
HWNEP0102I value Volumes found.	1609
HWNEP0103I Probing Configured Disks for Storage System: value.	1609
HWNEP0104I Number of Configured Disks Found Currently: value. Continuing to Probe Disks.	1609

HWNEP0105I value Configured Disks Found.	1609
HWNEP0106I Probing Views of Host Initiator access to Volumes.	1609
HWNEP0107I Finished probing Views.	1610
HWNEP0108I Initializing Probe for storage system storage system identification.	1610
HWNEP0109I Probe for storage system storage system identification completed.	1610
HWNEP0110I Collecting Nodes and fibre channel ports for storage system storage system identification.	1610
HWNEP0111E The connection to the storage device failed. The error code is error_code.	1610
HWNEP0113E The cluster IP address is not specified in the configuration file.	1611
HWNEP0114E The trustore location is not specified in the configuration file.	1611
HWNEP0115E The IBM Spectrum Control data is out of synch with the device configuration and a re-probe is required for device device name .	1611
HWNEP0116E The user configured for the subsystem subsystem name is not permitted to perform the requested action.	1611
HWNEP0117E The virtual disk (Vdisk)-to-host mapping was not created because the volume vdiskName is already mapped to the hostName host for the Device deviceName	1612
HWNEP0115I Starting Control Process for storage system storage system identification.	1612
HWNEP0116I Started deletion of volume VolumeID on subsystem Subsystem .	1612
HWNEP0117I Volume deletion completed successfully. Volume VolumeID on subsystem Subsystem was deleted.	1612
HWNEP0118I Started adding Managed Disk(s) Managed Disk ID to Managed-disk group Managed Disk group name on subsystem Subsystem.	1612
HWNEP0119I Finished adding Managed Disk(s) Managed Disk ID to Managed-disk group Managed Disk group name on subsystem Subsystem.	1613
HWNEP0120I Started creation of volume with size Size in pool Pool on subsystem Subsystem	1613
HWNEP0121I Volume creation completed successfully. New volume VolumeID created with size Size in pool Pool on subsystem Subsystem .	1613
HWNEP0122I Started assignment of volume VolumeID on subsystem Subsystem to initiator port Initiator Port on host Host .	1613
HWNEP0123I Finished assignment of volume VolumeID on subsystem Subsystem to initiator port Initiator Port on host Host Name .	1613
HWNEP0124I Started unassignment of volume VolumeID on subsystem Subsystem from initiator port Initiator Port on host Host Name .	1614
HWNEP0125I Finished unassignment of volume VolumeID on subsystem Subsystem from initiator port Initiator Port on host Host Name .	1614
HWNEP0126I Started creation of host host name on subsystem Subsystem with initiator ports WWPNs .	1614
HWNEP0127I Finished creation of host host name on subsystem Subsystem with initiator ports WWPNs .	1614
HWNEP0128I Host name hostName already exists for the WWPNs wwpns on the device Subsystem	1614
HWNEP0129E The operation failed because the device returned unexpected values.	1615
HWNEP0130E A IBM XIV CLI command failed. The error is error_code.	1615
HWNEP0131I The host definition for host host name on subsystem Subsystem contains additional Hostports WWPNs that will also be assigned to Volume VolumeID .	1615
HWNEP0132E The unassignment of Volume VolumeID from hostport WWPN failed because the definition for host host name on subsystem Subsystem contains additional hostports WWPNs .	1615
HWNEP0133E Error invoking the external process for device device name .	1615
HWNEP0134E Following exception occurred: exception .	1616
HWNEP0135E External process failed with error code error code .	1616
HWNEP0136E Error connecting to IP address with user ID user ID .	1616
HWNEP0137I Job job ID submitted for device device name .	1616
HWNEP0138I External process was successfully executed for device device name .	1616
HWNEP0139I An instruction was issued to add a copy of the volume_name volume_size-byte volume in the pool_name pool on the storage_system_name storage system.	1617
HWNEP0140I The copy of the volume_name volume_size-byte volume with the copy ID of VolumeID in the pool_name pool on the storage_system_name storage system was added successfully.	1617
HWNEP0141I Probing Internal Drives for Storage System: value.	1617
HWNEP0142I Number of Internal Drives Found Currently: value. Continuing to Probe Internal Drives.	1617
HWNEP0143I value Internal Drives Found.	1617
HWNEP0144I Probing Pools for Storage System: value.	1618
HWNEP0145I Number of Pools Found Currently: value. Continuing to Probe Pools.	1618
HWNEP0146I value Pools Found.	1618
HWNEP0147I Collecting asset and status information about storage_system_id storage system.	1618
HWNEP0148I Collecting cluster information for storage_system_id storage system.	1618
HWNEP0149I Collecting file system exports for storage_system_id storage system.	1618
HWNEP0150I Collecting nodes for storage_system_id storage system.	1619
HWNEP0151I Collecting file systems for storage_system_id storage system.	1619
HWNEP0152I Collecting pools for storage_system_id storage system.	1619
HWNEP0153I Collecting file system storage for storage_system_id storage system.	1619
HWNEP0154I Collecting filesets for storage_system_id storage system.	1619
HWNEP0155I Collecting links between file systems and nodes for storage_system_id storage system.	1620
HWNEP0156I Collecting quotas for storage_system_id storage system.	1620
HWNEP0157I Collecting file system snapshots for storage_system_id storage system.	1620
HWNEP0158I Collecting capacity for file_system_id file system.	1620
HWNEP0159I Creating the export export name on cluster cluster name .	1620
HWNEP0160I The export export name on cluster cluster name with path export path was created.	1620
HWNEP0161I The export export name on cluster cluster name is being changed.	1621
HWNEP0162I The export export name on cluster cluster name was changed.	1621
HWNEP0163I Setting quota quota type - quota name on file system file system name .	1621
HWNEP0164I Quota quota type - quota name on file system file system name has been created.	1621
HWNEP0165I Checking quota on file system file system name .	1621
HWNEP0166I Quota on file system file system name has been checked.	1621
HWNEP0167I The export export name on cluster cluster name is being removed.	1622
HWNEP0168I The export export name on cluster cluster name was removed.	1622
HWNEP0169E Command: command did not complete. IBM SONAS CLI message	1622

HWNEP0170I Creating fileset fileset name on file system files system name .	1622
HWNEP0171I Successfully created fileset fileset name on file system file system name .	1622
HWNEP0172I Removing fileset fileset name on file system files system name .	1622
HWNEP0173I Successfully removed fileset fileset name on file system file system name .	1623
HWNEP0174I Modifying fileset fileset name on file system files system name .	1623
HWNEP0175I Successfully modified fileset fileset name on file system file system name .	1623
HWNEP0176I Creating file system file system on cluster cluster name .	1623
HWNEP0177I Successfully created file system file system on cluster cluster name .	1623
HWNEP0178I Changing file system file system on cluster cluster name .	1624
HWNEP0179I Successfully changed file system file system on cluster cluster name .	1624
HWNEP0180I Removing file system file system on cluster cluster name .	1624
HWNEP0181I Successfully removed file system file system on cluster cluster name .	1624
HWNEP0182I Mounting file system file system .	1624
HWNEP0183I Successfully mounted file system file system .	1625
HWNEP0184I Unmounting file system file system .	1625
HWNEP0185I Successfully unmounted file system file system .	1625
HWNEP0186I Linking fileset fileset on file system file system .	1625
HWNEP0187I Successfully linked fileset fileset on file system file system .	1625
HWNEP0188I Unlinking fileset fileset on file system file system .	1625
HWNEP0189I Successfully unlinked fileset fileset on file system file system .	1626
HWNEP0190E The IBM Spectrum Control server could not connect to IP address using the SSH protocol.	1626
HWNEP0191E The IBM Spectrum Control server could not authenticate with IP address using the SSH protocol.	1626
HWNEP0192E The IBM Spectrum Control server could not execute a command on the IBM Storwize V7000 Unified/IBM SONAS device at IP address .	1626
HWNEP0193E The command name command failed because the following command executed on the NAS device failed with the return code return code : command returned: command output	1626
HWNEP0195I modify fileset	1627
HWNEP0196I change export	1627
HWNEP0197I create export	1627
HWNEP0198I remove export	1627
HWNEP0199I create fileset	1627
HWNEP0200I link fileset	1628
HWNEP0201I remove fileset	1628
HWNEP0202I unlink fileset	1628
HWNEP0203I change filesystem	1628
HWNEP0204I create filesystem	1628
HWNEP0205I mount filesystem	1628
HWNEP0206I remove filesystem	1629
HWNEP0207I unmount filesystem	1629
HWNEP0208I check quota	1629
HWNEP0209I set quota	1629
HWNEP0210I probe	1629
HWNEP0211W The command name command completed, however during post-processing the following command executed on the NAS device failed with the return code return code : command returned: command output As a result, the IBM Spectrum Control database is now out of sync with the current state of the NAS device.	1629
HWNEP0212I create disk in modifying file system	1630
HWNEP0213I Started deletion of host host name on subsystem Subsystem .	1630
HWNEP0214I Finished deletion of host host name on subsystem Subsystem .	1630
HWNEP0215I Collecting cache information for storage_system_id storage system.	1630
HWNEP0216I remove cached source	1630
HWNEP0217I create cached node	1631
HWNEP0218I remove cached node	1631
HWNEP0219I create cache	1631
HWNEP0220I remove cache	1631
HWNEP0221I modify cache source	1631
HWNEP0222I Creating cache source cache_source_name on cluster file_system_name.	1632
HWNEP0223I Created cache source cache_source_name on cluster file_system_name.	1632
HWNEP0224I Removing cache source cache_source_name on cluster file_system_name.	1632
HWNEP0225I Removed cache source cache_source_name on cluster file_system_name.	1632
HWNEP0226I Modifying cache source cache_source_name on cluster file_system_name.	1632
HWNEP0227I Modified cache source cache_source_name on cluster file_system_name.	1633
HWNEP0228I Creating cache cache_name on file system file_system_name.	1633
HWNEP0229I Created cache cache_name on file system file_system_name.	1633
HWNEP0230I Removing cache cache_name on file system file_system_name.	1633
HWNEP0231I Removed cache cache_name on file system file_system_name.	1633
HWNEP0232I Modifying cache cache_name on file system file_system_name.	1634
HWNEP0233I Modified cache cache_name on file system file_system_name.	1634
HWNEP0234I modify cache	1634
HWNEP0235I create cached source	1634
HWNEP0236I Configuring nodes node_names as cached nodes.	1634
HWNEP0237I Configured nodes node_names as cached nodes.	1634
HWNEP0238I Unconfiguring cached nodes node_names.	1635
HWNEP0239I Unconfigured cached nodes node_names.	1635

HWNEP0240I Executed control operation on cache cache_name on filesystem file_system_name .	1635
HWNEP0241I control cache	1635
HWNEP0242I run prepop	1635
HWNEP0243I list prepop	1636
HWNEP0244I Retrieving cache prepopulation status for file system file system name .	1636
HWNEP0245I Cache prepopulation status for file system file system name has been retrieved.	1636
HWNEP0246I Prepopulate cache data for fileset fileset_name on file system file_system_name using policy policy_name.	1636
HWNEP0247I Command to pre populate cached data for fileset fileset_name was successful.	1636
HWNEP0248W An error was encountered while parsing protocol options for export export_name. The options were not persisted, the probe will continue.	1637
HWNEP0249W The connection to the storage device failed. The error code is error_code.	1637
HWNEP0250I Started adding initiator port(s) initiator ports to host host name on subsystem subsystem .	1637
HWNEP0251I Finished adding initiator port(s) initiator ports to host host name on subsystem subsystem .	1637
HWNEP0252W A CLI command completed with warning. The warning message is : warning_message	1638
HWNEP0253W Volume creation completed with warning. New volume VolumeID created with size Size in pool Pool on subsystem Subsystem .	1638
HWNEP0254W Volume deletion completed with warning. Volume VolumeID on subsystem Subsystem was deleted.	1638
HWNEP0255I The task to execute the recommendations for optimizing the volumes on the storage system with an ID of storage_system_id was paused.	1638
HWNEP0256I The task for optimizing the volumes on the storage system with an ID of storage_system_id was canceled.	1638
HWNEP0257I The task for optimizing the volumes on the storage system with an ID of storage_system_id was resumed.	1639
HWNEP0258E The optimization task cannot be paused because the synchronization rate for the volume cannot be reset. The ID of the volume is volume_id and the ID of the storage system is storage_system_id.	1639
HWNEP0259E The optimization task cannot be resumed because the synchronization rate for the volume cannot be reset. The ID of the volume is volume_id and the ID of the storage system is storage_system_id.	1639
HWNEP0260I Started creation of host port host port name on storage system Storage System with initiator port WWPN .	1639
HWNEP0261I Finished creation of host port host port name on storage system Storage System with initiator port WWPN .	1639
HWNEP0262E The recommendation for the volume_name volume was not implemented because the command that was issued by the storage virtualizer returned the following error: error_message	1640
HWNEP0263I The synchronization of the volume_name volume with the volume copy was successful.	1640
HWNEP0264E The synchronization of the volume_name volume with the volume copy was unsuccessful.	1640
HWNEP0265E The CLI command that was issued for the storage_system_name storage system failed and generated the following error: error_message	1640
HWNEP0266I Started expanding the capacity of volume volume on subsystem subsystem from oldsize to newsize bytes.	1641
HWNEP0267I Finished expanding the capacity of volume volume on subsystem subsystem to newsize bytes.	1641
HWNEP0268E The server operating system or version is not supported by IBM Spectrum Control for IBM Spectrum Scale.	1641
HWNEP0269E The IBM Spectrum Scale cluster information cannot be displayed. All the nodes in the cluster are down or cannot be contacted.	1641
HWNEP0270E The switch cannot respond to SNMP queries because of an authentication error.	1641
HWNEP0271E The following password decryption exception occurred: exception	1642
HWNEP0272E The switch cannot respond to SNMP queries because of the following exception: exception	1642
HWNEP0273E The following exception occurred because the OID format is incorrect: exception	1642
HWNEP0274E The switch cannot respond to SNMP queries because of a timeout problem.	1642
HWNEP0270I Retrieved the file module address file_module_address.	1643
HWNEP0271I No quota data was collected. Quota limits are not activated for the file systems that are associated with the IBM Spectrum Scale cluster.	1643
HWNEP0272I Collecting file systems that are mounted on the nodes of storage system storage_system_id.	1643
HWNEP0275W One or more operations failed for the CLI command that was issued for the storage system. The following error was generated: errorMsg .	1643
HWNEP0276E Command execution failed because sudo is not installed.	1643
HWNEP0277I Commands are executed through 'sudo'.	1644
HWNEP0278E User can not execute command through sudo.	1644
HWNEP0279I Collecting remote file systems for storage_system_id storage system.	1644
HWNEP0280I Collecting remote file systems that are mounted on the nodes of storage system storage_system_id.	1644
HWNEP0281E The switch is returning corrupted data.	1644
HWNEP0282E Zoning data cannot be collected because there is a transaction in progress on the switch	1644
HWNEP0283E VSAN vsan_name was not found.	1645
HWNEP0284E No zoning data collected from the switch.	1645
HWNEP0285E Cannot authenticate to the object storage using the specified user credentials.	1645
HWNEP0286E An object storage request failed on the GPFS cluster.	1645
HWNEP0287E Error when collecting Accounts information from Object Storage Service using REST protocol.	1645
HWNEP0288E Error when collecting Containers information from Object Storage Service using REST protocol.	1646
HWNEP0281I Collecting object storage accounts for storage_system_id storage system.	1646
HWNEP0282I Collecting object storage containers for storage_system_id storage system.	1646
HWNEP0289E Failed to retrieve container information because the number of containers now exceeds the maximum number of containers that can currently be collected for an account (MAX Containers).	1646
HWNEP0290E The probe failed to retrieve object storage account information from the storage system storage_system_id because the userid user does not have the required authority.	1647
HWNEP0291E The probe failed to retrieve object storage container information from the storage system storage_system_id because the userid user does not have the required authority.	1647
HWNEP0292E Cannot query the object service for information about accounts and containers as the specified user does not have admin privileges.	1647
HWNEP0293W The probe did not collect information about all the object accounts for the storage system storage_system_id as the userid user does not have sufficient authority on the storage system.	1647
HWNEP0294W An authentication error prevented the switch from responding to SNMP queries regarding the ability of the switch to perform zone control.	1648
HWNEP0295W A timeout prevented the switch from responding to SNMP queries regarding the ability of the switch to perform zone control.	1648
HWNEP0296W The switch cannot respond to SNMP queries to check the ability of the switch to perform zone control because of the following exception: exception	1648
HWNEP0297W The switch cannot respond to SNMP queries to check the ability of the switch to perform zone control because of the following exception: exception	1649
HWNEP0298I Collecting IBM Cloud Object Storage configuration.	1649

HWNEP0299I Collecting IBM Cloud Object Storage vaults.	1649
HWNEP0300I Collecting detailed IBM Cloud Object Storage status.	1649
HWNEP0301W The IP address ip_address for the FlashSystem storage system is not the management IP address.	1649
HWNEP0302I Collecting Transparent Cloud Tiering information for storage_system_id storage system.	1650
HWNEP0303I No Transparent Cloud Tiering configuration was detected on the IBM Spectrum Scale cluster.	1650
HWNEP0304E Cannot connect to IBM Cloud Object Storage.	1650
HWNEP0305I Collecting disk controllers for storage system storage system identification.	1650
HWNEP0306I Collecting disks for storage system storage system identification.	1650
HWNEP0307I Collecting CIFS shares for storage system storage system identification.	1651
HWNEP0308I Collecting NFS exports for storage system storage system identification.	1651
HWNEP0309I The data is being collected by the data collector: data collector host.	1651
HWNEP0310I Discovery found number storage systems.	1651
HWNEP0311I Probing nodes or directors for storage system name storage system.	1651
HWNEP0312I Probe found number nodes or directors.	1652
HWNEP0313I Probing pools for storage system name storage system.	1652
HWNEP0314I Probe found number pools.	1652
HWNEP0315I Probing disk groups for storage system name storage system.	1652
HWNEP0316I Probe found number disk groups.	1652
HWNEP0317I Probing disks for storage system name storage system.	1652
HWNEP0318I Probe found number disks.	1653
HWNEP0319I Probing host connections for storage system name storage system.	1653
HWNEP0320I Probing ports for storage system name storage system.	1653
HWNEP0321I Probing volumes for storage system name storage system.	1653
HWNEP0322I Probe found number volumes. Continuing to probe volumes.	1653
HWNEP0323I Probe found number volumes for storage system name storage system.	1654
HWNEP0324I Probing NAS nodes for storage system name storage system.	1654
HWNEP0325I Probe found number NAS nodes.	1654
HWNEP0326I Probing file systems that are mounted on the NAS nodes of storage system name storage system.	1654
HWNEP0327I Probe found number file systems.	1654
HWNEP0328I Probing file system exports for storage system name storage system.	1654
HWNEP0329W profile name version version number SMI-S Profile is not supported.	1655
HWNEP0330E Unable to find minimum required SMI-S profile to proceed with requested task.	1655
HWNEP0331I Probing copy pair relationships for storage system name storage system.	1655
HWNEP0332I Probe found number copy pairs.	1655
HWNEP1111E There is no connection for the specified device.	1655
HWNEP1112E No SSH server found on the device.	1656
HWNEP1113E Unsupported version.	1656
HWNEP1114E The connection to the device failed.	1656
HWNEP1115E Authentication failed.	1656
HWNEP1116E Unknown host.	1656
HWNEP1117E The passphrase is wrong.	1657
HWNEP1118E The passphrase is missing.	1657
HWNEP1119E Unknown error.	1657
HWNEP1120E ESSNI not available.	1657
HWNEP1121E Private key not found.	1657
HWNEP1122E Invalid format for the private key.	1657
HWNEP1123E Unable to establish a connection to the device through http port 80.	1658
HWNEP1124I Log collection successfully started for storage system name storage system.	1658
HWNEP1125E The activity requested is already in progress on storage system name storage system.	1658
HWNEP1126I The support log activity has started successfully storage system name storage system.	1658
HWNEP0112E The CLI command that was issued for the storage system failed and generated the following error: error_message	1658
HWNEP1127I The probe failed to retrieve encryption information from the storage system storage_system_id because the userid user does not have the required authority.	1659
HWNEP1128E The process failed because it was unable to find the Export Tool. Expected location was loc of tool.	1659
HWNEP1129E The process failed because the userid or password provided failed to connect to the Export Tool.	1659
HWNEP1130E The process failed because the Hitachi SVP was busy and did not return data or timed out.	1659
HWNPM5412E Performance statistics collection is not enabled.	1659
HWNEP1131E The process failed because the Hitachi performance interval is set to something other than 1 or 5 minutes.	1660
HWNEP1132W Can't collect further system information because the device returned unexpected values.	1660
NAD0001I Connecting to hostname using protocol protocol.	1660
NAD0002W Connection to hostname failed using protocol protocol: error.	1660
NAD0003I Connected to hostname using protocol protocol.	1660
NAD0005E Connection to hostname failed using protocol protocol: error message.	1660
NAD0006E Exception thrown for method method name: error message.	1661
NAD0007I Closing connection to hostname.	1661
NAD0008E Invalid protocol protocol passed to method name.	1661
NAD0010E Invalid parameter(s) parameter name passed to method name.	1661
NAD0013I Installing GUID on remote machine: hostname.	1661
NAD0014I GUID successfully installed on remote machine: hostname.	1661
NAD0018E Command on remote machine: host name failed. Error code = value executing command value.	1661
NAD0019E Parameter parameter passed to method is null or 0 length.	1661
NAD0055E Failed to connect to remote host host.	1661

NAD0097I Opening connection to hostname.	1662
NAD0180I Installing re-distributable package on .	1662
NAD0181I Install of re-distributable package on succeeded.	1662
NAD0182E Failed to install re-distributable package on .	1662
NAD0186I Trying to locate package TIVguid using pkginfo ...	1662
NAD0187I Package TIVguid is not installed.	1662
NAD0188I Checking TIVguid default install path : path ...	1662
NAD0259W Unable to determine Storage Resource Agent version on host . Fabric Discovery will not be invoked.	1662
NAD0145E Cannot get version information from agent on host .	1663
NAD0146E The connection to remote machine failed because the Remote Execution and Access component was unable to create a temporary directory on the remote machine. Remove unneeded ~CSRI* directories in the remote machine's temporary directory.	1663
NAD0156E The server host_address cannot be reached because the host name or IP address is not recognized.	1663
NAD0157E The server host_name cannot be contacted. The server might be down, unreachable due to network problems, or the SSH credentials might be invalid.	1663
NAD0260I Agent is active.	1664
NAD0272W The connection to the Storage Resource Agent on host name was not established. Retrying using the IP address.	1664
NAD0274E An SSH certificate certificate name already exist.	1664
NAD0275E Failed to connect to remote host hostname and port. Failed to establish a secure connection.	1664
NAD0276E Failed to connect to remote host hostname and port. Failed to establish a secure connection because the SSL handshake failed.	1664
NAD0277E Failed to connect to remote host hostname and port. Failed to establish a secure connection because of an invalid SSL key.	1664
NAD0278E Failed to connect to remote host hostname and port. Failed to establish a secure connection because the identity of the peer could not be verified.	1665
NAD0279E Failed to connect to remote host hostname and port. Failed to establish a secure connection because of an SSL protocol error.	1665
NAD0281E The Storage Resource agent cannot be deployed because of insufficient space or other issues on the target system. The error is: error message.	1665
BTAVM2272W Unsupported virtual disk backing info for disk "Disk name" of hypervisor Hypervisor name, virtual machine "VM name": Virtual disk type.	1665
BTAVM2273W Unable to find file "File name" which is the backing device of the virtual disk "Disk name" of hypervisor Hypervisor name, virtual machine "VM name".	1666
BTAVM2274W Probe of hypervisor Name of the Hypervisor completed with warnings.	1666
BTAEC - Event correlator messages	1666
BTAEC1020W The Device Server cannot listen for Forwarded SNMP Traps on port {0}. Port {1} will be used instead.	1666
BTAHM - Host manager messages	1666
BTAHM2501E The service name Service failed to start due to condition.	1667
BTAHM2520E Agent agent name has been marked inactive.	1667
BTAHM2521E The agent returned an invalid name.	1667
BTAHM2522E agent name is not a known agent.	1667
BTAHM2524E The agent returned an invalid port number.	1667
BTAHM2525E Agent agent name cannot be removed because it is active.	1667
BTAHM2527E Unexpected error java error.	1668
BTAHM2528I Agent host name has been marked active.	1668
BTAHM2551I An inactive agent agent name has been removed.	1668
BTAHM2580I The service name service started.	1668
BTAHM2581I The service name service is shut down.	1668
BTAIC - Inband change agent messages	1668
BTAIC1200E The InbandChangeAgent cannot contact the EventCorrelator.	1668
BTAIC1201E An error occurred while reading the InbandEvents file.	1669
BTAIC1202E The InbandChangeAgent thread has been interrupted.	1669
BTAIC1203E The InbandChangeAgent failed to execute the Event.exe command.	1669
BTAIC1204E The AIX protocol driver must be uninstalled to prevent it from interfering with the EventScanner.	1669
BTAIC1205E In-band event notification requires at least maintenance level 2 for AIX 5.2.	1669
BTAIC1206E In-band event notification requires at least maintenance level 1 for AIX 5.3.	1670
BTAIC1207E The version of AIX that is running on this managed host is not supported.	1670
BTALG - Logging toolkit messages	1670
BTALG0001I Logging Toolkit is ready.	1672
BTALG0002I log add <logger_name>	1672
BTALG0003I -handler=<handler_name>	1672
BTALG0004I log debug {on off}	1672
BTALG0005I log get	1673
BTALG0006I -filterkey	1673
BTALG0007I -locale	1673
BTALG0008I -format	1673
BTALG0009I -maxfiles	1673
BTALG0010I -maxfilesize	1673
BTALG0011I log get <object_name>	1673
BTALG0012I -filename	1673
BTALG0013I -filterkey	1673
BTALG0014I -formatter	1674
BTALG0015I -handlers	1674
BTALG0016I -locale	1674
BTALG0017I -logstate	1674
BTALG0018I -maxfiles	1674
BTALG0019I -maxfilesize	1674
BTALG0020I log help	1674
BTALG0021I log list	1674
BTALG0022I -formatters	1674

BTALG0023I -locales	1675
BTALG0024I -loggers	1675
BTALG0025I -handlers	1675
BTALG0026I log remove <logger_name>	1675
BTALG0027I -handler=<handler_name>	1675
BTALG0028I log set	1675
BTALG0029I -defaults	1675
BTALG0030I -filterkey <INFO ERROR WARN>	1675
BTALG0031I -format {plain_text pdxml}	1676
BTALG0032I -locale {<locale> default}	1676
BTALG0033I -maxfiles <max_files>	1676
BTALG0034I -maxfilesize <max_file_size>	1676
BTALG0035I log set <object_name>	1676
BTALG0036I -filename <file_name>	1676
BTALG0037I -filterkey <INFO ERROR WARN>	1676
BTALG0038I -formatter <formatter_name>	1676
BTALG0039I -locale <locale>	1676
BTALG0040I -logstate {on off}	1677
BTALG0041I -maxfiles <max_files>	1677
BTALG0042I -maxfilesize <max_file_size>	1677
BTALG0043I Invalid number of parameters.	1677
BTALG0044I Invalid option.	1677
BTALG0045I Function not supported for native loggers.	1677
BTALG0046I Locale is set to locale.	1677
BTALG0047I Logger is state.	1677
BTALG0048I on	1677
BTALG0049I off	1678
BTALG0050I Attached handlers are handlers.	1678
BTALG0051I Filterkey is set to filterkey.	1678
BTALG0052I Format is set to format.	1678
BTALG0053I Formatter is set to formatter.	1678
BTALG0054I Filename is set to filename.	1678
BTALG0055I Maxfiles is set to maxfiles.	1678
BTALG0056I Maxfilesize is set to maxfilesize KB.	1678
BTALG0057I Locale was set to locale.	1679
BTALG0058I Filterkey was set to filterkey.	1679
BTALG0059I Invalid format format.	1679
BTALG0060I Format was set to format.	1679
BTALG0061I Formatter was set to format.	1679
BTALG0062I Filename was set to filename.	1679
BTALG0063I Maxfiles was set to maxfiles.	1679
BTALG0064I Maxfilesize was set to maxfilesize KB.	1679
BTALG0065I Logging defaults have been restored.	1679
BTALG0066I Failed to update property.	1680
BTALG0067I Logger has been turned state.	1680
BTALG0068I Invalid option option.	1680
BTALG0069I -handler	1680
BTALG0070I Handler added successfully.	1680
BTALG0071I Failed to add handler.	1680
BTALG0072I Handler removed successfully.	1680
BTALG0073I Failed to remove handler.	1680
BTALG0074I Invalid command command.	1680
BTALG0075I Debug is set to state.	1681
BTALG0076I plain_text	1681
BTALG0077I pdxml	1681
BTALG0078I Failed to get property information.	1681
BTALG0079I Displays logging properties.	1681
BTALG0080I Provides general information on the Logging Service commands.	1681
BTALG0081I Defines logging properties.	1681
BTALG0082I Adds a handler to the specified logger.	1681
BTALG0083I Enables or disables additional logging commands.	1682
BTALG0084I Provides a list of loggers, handlers, or formatters.	1682
BTALG0085I Removes a handler object.	1682
BTALG0086I log add <logger_name> [option]	1682
BTALG0087I log debug {on off}	1682
BTALG0088I log get [option]	1682
BTALG0089I log help [option]	1682
BTALG0090I log list [option]	1682
BTALG0091I log remove <logger_name> [option]	1682
BTALG0092I log set [option]	1683
BTALG0093I IBM Spectrum Control Logging Toolkit for Fabric	1683
BTALG0094I Command Line Interface - Version version Release release Level level minor	1683

BTALG0095I LOGGING SERVICE COMMANDS	1683
BTALG0097I Command	1683
BTALG0098I Description	1683
BTALG0099I See	1683
BTALG0100I -add	1683
BTALG0101I -debug	1684
BTALG0102I -get	1684
BTALG0103I -help	1684
BTALG0104I -list	1684
BTALG0105I -remove	1684
BTALG0106I -set	1684
BTALG0107I OPTION	1684
BTALG0108I COMMAND	1684
BTALG0109I DESCRIPTION	1684
BTALG0110I add	1685
BTALG0111I debug	1685
BTALG0112I get	1685
BTALG0113I list	1685
BTALG0114I remove	1685
BTALG0115I set	1685
BTALG0116I Adds the handler to the specified logger.	1685
BTALG0117I log get <object_name> [option]	1685
BTALG0118I Displays the current types of messages that are logged in the log file.	1686
BTALG0119I Displays the maximum number of log files to be created.	1686
BTALG0120I Displays the maximum file size of the log before a new log file is created.	1686
BTALG0121I Displays the format in which messages are saved in the log file.	1686
BTALG0122I Displays the current language locale setting in which messages are displayed in the message log file.	1686
BTALG0123I Displays the file name associated with the specified handler.	1686
BTALG0124I Displays the formatter that is attached to the specified handler.	1686
BTALG0125I Displays the handler that is attached to the specified logger.	1686
BTALG0126I Displays if the logger is on or off. You must specify a logger for the object name.	1687
BTALG0127I Displays the current types of messages that are logged for the specified logger.	1687
BTALG0128I Displays the maximum number of log files to be created for the specified handler.	1687
BTALG0129I Displays the maximum file size of log files created by the specified handler.	1687
BTALG0130I log set <object_name> [option]	1687
BTALG0131I Logging configuration corrupted. Restoring default configuration.	1687
BTALG0132I Displays a list of loggers.	1687
BTALG0133I Displays a list of formatters.	1688
BTALG0134I Displays a list of handlers.	1688
BTALG0135I Removes the handler from the specified logger.	1688
BTALG0136I -handlers=<list_of_handlers>	1688
BTALG0137I Specifies the types of messages that will be logged.	1688
BTALG0138I Sets the maximum number of log files to be created.	1688
BTALG0139I Sets the maximum file size (in kilobytes) of the log before a new log file is created.	1688
BTALG0140I Sets the format in which messages are saved in the log file.	1688
BTALG0142I Sets the file name where the specified handler will output log messages.	1689
BTALG0143I Sets the formatter used by the specified handler.	1689
BTALG0145I Turns the log on or off. You must specify a logger for the object name.	1689
BTALG0146I Specifies the types of messages that will be logged.	1689
BTALG0147I Sets the maximum number of log files to be created.	1689
BTALG0148I Sets the maximum file size (in kilobytes) of the log before a new log file is created.	1689
BTALG0149I Resets the logging properties to their default settings.	1689
BTALG0150I help	1689
BTALG0151I State	1690
BTALG0152I Filter	1690
BTALG0153I Handlers	1690
BTALG0154I Logger	1690
BTALG0155E logger is not a valid logger. Failed to update property.	1690
BTAMS - Spectrum Control Messaging Service messages	1690
BTAMS0500I IBM Spectrum Control MessagingService started successfully.	1690
BTAMS0501I The Messaging Service has shutdown.	1690
BTAMS0502I Service service name subscribed to topic topic name.	1691
BTAMS0503I Event published to topic topic name.	1691
BTAMS0504E Messaging Service failed to get a proxy to the service name service.	1691
BTAMS0505E Messaging Service could not invoke the onMessage method on service service name.	1691
BTAMS0001W Failed to load the configuration for the database exception handler.	1691
BTAMS0002I Data server	1691
BTAMS0003I Device server	1692
BTAMS0004I Unknown	1692
BTAMS0005W Failed to update the database pool monitor handler handler name.	1692
BTAMS0006W Failed to initialize the database pool monitor handler handler name.	1692
BTAQE - Spectrum Control Query Engine messages	1692

BTAQE1100E Query Engine Event Generator can not start.	1693
BTAQE1101E Unable to open the database.	1693
BTAQE1102E Unable to close the database.	1693
BTAQE1104E The Query Engine check write authority failed.	1693
BTAQE1105E Check for QueryEngine Authentication failed.	1694
BTAQE1106E The SANQueryEngine thread has been interrupted.	1694
BTAQE1109E An error occurred while attempting to save the IP target to the database.	1694
BTAQE1110E An error occurred while attempting to delete an IP target from database.	1694
BTAQE1111E An error occurred while querying the IP target information from the database.	1695
BTAQE1116E Database errors occurred while performing queries on Tasks.	1695
BTAQE1117E Database errors occurred while saving the task.	1695
BTAQE1118E Errors occurred while resolving InterconnectElement and Port relationship.	1695
BTAQE1119E Errors in Topology XML generator.	1696
BTAQE1120E Errors in creating an entity.	1696
BTAQE1121E Invalid target host IP address.	1696
BTAQE1122E Failed to get the SNMP Service proxy.	1696
BTAQE1123E Unable to close the event publisher.	1697
BTAQE1124E Unable to compress scanner result due to IOException: exception.	1697
BTAQE1125E Unable to uncompress the scanner result due to IOException: exception.	1697
BTAQE1126I An unidentified port was removed from the scan data.	1697
BTAQE1127E An outband scanner failed to capture the scan data.	1697
BTAQE1128E An outband scanner failed to save the scan data for benchmark comparison.	1698
BTAQE1129E An outband scanner failed to read the benchmark file benchmark file name saved from the previous scan.	1698
BTAQE1130E An outband scanner failed to decrypt the password for target target IP.	1698
BTAQE1134I The outband agent target address TargetIP address does not respond to Fibre Channel MIB (previously called the Fibre Alliance MIB) queries.	1698
BTAQE1135E Unable to get the license state from the license server.	1699
BTAQE1136E The Query Engine cannot obtain a valid IP address for the host target.	1699
BTAQE1137E The Query Engine could not obtain the information for target host target from the database.	1699
BTAQE1138E The Query Engine could not obtain the capability information for target host target from the database.	1699
BTAQE1139E The Query Engine could not obtain the information for all known target hosts from the database.	1700
BTAQE1140E Error creating an event subscriber.	1700
BTAQE1141E The Query Engine could not obtain the scheduled scan information from the database.	1700
BTAQE1142E The Query Engine could not obtain the list of active scanners from the database.	1700
BTAQE1143E The Query Engine could not obtain the list of inactive scanners from the database.	1701
BTAQE1144E An error occurred attempting to run the scanner name scanner on the IBM Spectrum Control managed host target.	1701
BTAQE1145E The scanner name scanner running on the IBM Spectrum Control managed host target found no SAN.	1701
BTAQE1146E The scanner name scanner running on IBM Spectrum Control managed host target found no host-based adapter (HBA).	1701
BTAQE1147E The scanner name scanner running on IBM Spectrum Control managed host target found no SCSI host-based adapter (HBA).	1701
BTAQE1149E A scanner overlap condition has occurred for the scanner name scanner on the IBM Spectrum Control managed host target .	1702
BTAQE1150I The outband agent target address TargetIP address does not support topology discovery through SNMP Fibre Channel MIB (previously called the Fibre Alliance MIB), or Cisco VSAN MIB queries.	1702
BTAQE1151I The outband agent with target address TargetIP address has been added.	1702
BTAQE1152I The outband agent with target address TargetIP address has been removed.	1702
BTASA - Spectrum Control SAN scanner agent messages	1703
BTASA1400E The SAN Agent Scanner failed to execute the inband scanner scanner with the command: scan command.	1703
BTASA1401E The SAN Agent Scanner failed to capture the inband scan data.	1703
BTASA1403E The SAN Agent Scanner failed to save the scan data for benchmark comparison.	1703
BTASA1404E The SAN Agent Scanner failed to read the benchmark file benchmark file saved from the previous scan.	1704
BTASA1405E The SAN Agent Scanner failed to retrieve the global unique identifier from Host Query.	1704
BTASA1406I The SAN Agent Scanner Service has initialized successfully.	1704
BTASA1407I The Inband scanner scanner has started.	1704
BTASA1408I The Inband scanner scanner has ended with return code return code.	1704
BTASA1409E A scanner overlap condition has occurred on the IBM Spectrum Control managed host.	1705
BTASA1420E The GS-3 Zone Control DLL could not be loaded.	1705
BTASD - Fabric User Interface messages	1705
BTASD1922E An error occurred while getting the information from device services.	1706
BTASD1923E The agent agent identifier is currently agent state. It must be in order to remove it.	1706
BTASD1930E Unable to contact zone agent	1706
BTASD1931E Unable to contact zoning agent. Token used for contacting zone agent is invalid.	1706
BTASD1932E Agent capable of configuring zoning could not be found on this Fabric	1707
BTASD1933E Zoning is already being configured on this Fabric. New zoning can not be done until agent is available again	1707
BTASD1934E Unable to delete the selected entities.	1707
BTASD1935E The delete failed because an agent is still installed on the selected computer or configured in Data Agent or Inband Fabric Agent list.	1707
BTASD1936E Unexpected server response message_status= agent response.	1708
BTASD1937E A zone name cannot contain the characters '!', '%', '*' or '!' in its name. Brocade zone names also cannot contain '\$' nor '-'. The first character for a zone name must be alphanumeric. Enter a new name for this zone.	1708
BTASD1938E A zone set name cannot contain the characters '!', '%', '*' or '!' in its name. Brocade zone set names also cannot contain '\$' nor '-'. The first character for a zone set name must be alphanumeric. Enter a new name for this zone set.	1708
BTASD1939E Zoning is already being configured by lock owner on this Fabric since lock time. New zoning can not be done until agent is available again.	1708
BTASD1940E Zoning is already being configured by lock owner on this Fabric since lock time. New zoning can not be done until agent is available again. Do you want to release the lock from user {0}?	1708
BTASD1941E The lock for the zone control operations has been reset. Do you want to re-obtain the lock and continue zone operations?	1709

BTASD1942E Zoning changes cannot be made at this time. Zoning for this fabric is currently locked by lock owner since date.	1709
BTASD1943E An alias with the same name already exists. Enter a new name for this alias.	1709
BTASD1944E The Alias name field must be filled in to create an alias.	1709
BTASD1945E A alias name cannot contain the characters ':' ','%' or '!' in its name. Brocade alias names also cannot contain '\$' nor '-'. The first character for an alias name must be alphanumeric. Enter a new name for this alias.	1710
BTASD1946E An alias name cannot begin with a number. Enter a new name for this alias.	1710
BTASD1947E An alias, zone and/or zone set in the same configuration can not have the same name. Enter a new name.	1710
BTASD1948E This alias does not contain any members. Add a member to this alias.	1710
BTASD1949E There are too many members selected. The maximum number of members allowed for this alias is max members. Member(s) must be removed before you can continue.	1710
BTASD1950E The zone configuration has the maximum number of aliases allowed. The maximum number of aliases is max aliases. An existing alias must be deleted before a new one can be created.	1711
BTASD1951E Fabric fabric name has zone count zones with non-standard members: zone names. Zone and ZoneSet changes cannot be applied to zones with nonstandard members.	1711
BTASD1952E Zone set zoneset name is active or in activation or deactivation pending status. Deletion of the zone set is not supported for this fabric. After committing the deactivation, this zone set can be deleted in the next zone control session.	1711
BTASD1953E Renaming of the active zone set is not supported for this fabric. Deactivate zone set zoneset name before attempting to rename it.	1711
BTASD1954E Zoning cannot be done because there is no connection to the SMI-S provider. Reason: reason	1711
BTASD2001W Zoning has changed on this fabric since the configuration panel was opened. You may need to run the Fabric discovery/probe again to get the zone information within IBM Spectrum Control synchronized with the fabric. Do you still want to make zoning changes?	1712
BTASD2002W This alias is not assigned to any zones. This could result in an error when the zone configuration is applied later. Do you want to continue?.	1712
BTASD2003W No exclusive fabric-wide lock is available on the switches for fabric fabric name. Other users might be modifying the zoning configuration from outside of IBM Spectrum Control during your zoning operations.	1712
BTASD2004W This Out Of Band Agent is already defined with the same parameters. Would you like to save it anyway?	1712
BTASD3001I A probe job for fabric fabric name has been submitted. The inactive Zone Definition for this fabric will have the old Zone Definition until the probe job is complete. Wait a few minutes before working with Zone Definition for this fabric.	1713
BTATG - UNIX Command Line Interface (CLI) help messages	1713
BTATG0000E You must have root user authority to run this program.	1713
BTATG0001E Invalid option '&1'.	1713
BTATG0003E The format of the GUID is invalid.	1713
BTATG0004I A GUID already exists on this host. A new GUID will not be created.	1714
BTATG0005I A GUID entry was not found. The program is generating a new one.	1714
BTATG0006I A GUID entry was not found.	1714
BTATG0007E A GUID entry was not created.	1714
BTATG0008E The GUID entry could not be written.	1714
BTATG0009E The GUID entry can not be read.	1714
BTATG0011E When using '&1' you must enter '&3' or use '&3'.	1715
BTATG0012E The tivguid program encountered an internal error.	1715
BTATG0013E '&1' return status is '&2'.	1715
BTAVM	1715
BTVMW4001I Connection test to VMware VI Data Source VMware VI Data Source host name FAILED due to VMWareConnectionStatus	1715
BTVMW2013E The addition of the data source Name of the data source failed.	1716
BTVMW2014W This VMWare VI Data Source is already defined with the same parameters. Would you like to save it anyway?	1716
BTAZC - Zone control agent messages	1716
BTAZC0001E An error occurred while parsing the Zoning Configuration XML for SAN SAN_name.	1718
BTAZC0002E Failed to create or update zone set zone_set_name on the GS3 SAN SAN_name.	1719
BTAZC0003E Failed to delete zone set zone_set_name on the GS3 SAN SAN_name.	1719
BTAZC0004E Unable to start transaction on the SAN SAN_name.	1719
BTAZC0005E Unable to commit a transaction on the Brocade SAN SAN_name.	1719
BTAZC0006E Failed to deactivate the zone set zone_set_name on SAN SAN_name.	1720
BTAZC0007E Failed to activate the zone set zone_set_name on SAN SAN_name.	1720
BTAZC0008E Failed to delete the zone set zone_set_name on SAN SAN_name.	1720
BTAZC0009E Failed to delete the zone alias zone_alias_name on SAN SAN_name.	1720
BTAZC0010E Failed to create the zone alias zone_alias_name on SAN SAN_name.	1721
BTAZC0011E Failed to delete the zone zone_name on SAN SAN_name.	1721
BTAZC0012E Failed to create the zone zone_name on SAN SAN_name.	1721
BTAZC0013E Failed to create the zone set zone_set_name on SAN SAN_name.	1722
BTAZC0014E Failed to deactivate the zone set zone_set_name on SAN SAN_name.	1722
BTAZC0015E Failed to add the zone zone_name to the zone set zone_set_name on SAN SAN_name.	1722
BTAZC0016E One or more nonstandard zone members present in current zoning configuration for SAN SAN_name. Due to that, the Zone Control Agent will not attempt to modify the zoning configuration for the SAN.	1722
BTAZC0017E Start transaction for Zone Control failed.	1723
BTAZC0018E Rollback error.	1723
BTAZC0019E An error occurred while releasing a resource.	1723
BTAZC0020E An error occurred while creating a zone set.	1723
BTAZC0021E An error occurred while creating a zone.	1724
BTAZC0022E An error occurred while creating an alias.	1724
BTAZC0023E An error occurred while creating a member.	1724
BTAZC0024E An error occurred while adding a zone to a zone set.	1724
BTAZC0025E An error occurred while adding a member to a zone.	1725
BTAZC0026E An error occurred while adding an alias to a zone.	1725
BTAZC0027E An error occurred while adding a member to an alias.	1725
BTAZC0028E An error occurred while removing a zone from zone set.	1725

BTAZC0029E An error occurred while removing an alias from a zone.	1726
BTAZC0030E An error occurred while removing a member from a zone.	1726
BTAZC0031E An error occurred while removing a member from an alias.	1726
BTAZC0032E An error occurred while deleting a zone member	1726
BTAZC0033E An error occurred while deleting a zone.	1727
BTAZC0034E An error occurred while deleting a zone set.	1727
BTAZC0035E An error occurred while deleting an alias.	1727
BTAZC0036E An error occurred while activating a zone set.	1727
BTAZC0037E An error occurred while deactivating a zone set.	1728
BTAZC0038E An error occurred while pinging the Zoning Agent.	1728
BTAZC0039E An error occurred during transaction commit action.	1728
BTAZC0040E An error occurred while closing a session.	1728
BTAZC0041E An error occurred while saving the zone information.	1729
BTAZC0042E An error occurred during the Get Capabilities command	1729
BTAZC0043E An error occurred sending the zone control command array.	1729
BTAZC0044E An error occurred while sending commands to Switch.	1729
BTAZC0045E An error occurred: n unsupported zone database.	1730
BTAZC0046E A native error occurred: invalid field length.	1730
BTAZC0047E A native error occurred: invalid number of members.	1730
BTAZC0048E A native error occurred: invalid arguments.	1730
BTAZC0049E A native error occurred: null fabric handle.	1731
BTAZC0050E An unknown error occurred during Zone control.	1731
BTAZC0051E An XML parse error occurred during Zone Control operations.	1731
BTAZC0052E Unable to create logical zone definition.	1731
BTAZC0053E An error occurred during Zone Control: library not opened.	1732
BTAZC0054E Non standard members in the current zone definition.	1732
BTAZC0055E A native error occurred: function not supported.	1732
BTAZC0056E An error occurred: not connected to the SAN.	1732
BTAZC0057E A native error occurred: invalid buffer index.	1733
BTAZC0058E A native error occurred during an HBA API call.	1733
BTAZC0059E A native error occurred: no memory available.	1733
BTAZC0060E A native error occurred: error loading the HBA API.	1733
BTAZC0270E An error occurred during an HBA API call.	1734
BTAZC0271E An error occurred during an HBA API call: not supported.	1734
BTAZC0272E An error occurred during an HBA API call: invalid handle.	1734
BTAZC0273E Bad argument with the HBA API.	1734
BTAZC0274E An error occurred during an HBA API call: illegal WWN.	1735
BTAZC0275E An error occurred during an HBA API call: illegal index.	1735
BTAZC0276E Larger buffer required with the HBA API.	1735
BTAZC0277E Information has changed since the last call to HBA_RefreshInformation.	1735
BTAZC0278E SCSI check condition reported with the HBA API.	1736
BTAZC0279E HBA error: adapter may be busy or reserved. Retry may be effective.	1736
BTAZC0280E HBA API request timed out. Retry may be effective.	1736
BTAZC0281E Referenced HBA has been removed or deactivated.	1736
BTAZC0282E Extended Link Service reject with the HBA API.	1737
BTAZC0283E An error occurred during an HBA API call: invalid LUN.	1737
BTAZC0284E An error occurred during an HBA API call: incompatible.	1737
BTAZC0285E Ambiguous WWN with HBA API call.	1737
BTAZC0286E Local bus error with the HBA API.	1738
BTAZC0287E Local target error with the HBA API.	1738
BTAZC0288E Local LUN error with the HBA API.	1738
BTAZC0289E HBA API error: local SCSI bound.	1738
BTAZC0290E HBA API error on FCID target.	1739
BTAZC0291E Target node WWN error with the HBA API.	1739
BTAZC0292E Target port WWN error with the HBA API.	1739
BTAZC0293E Target LUN error with the HBA API.	1739
BTAZC0294E Target LUN ID error with the HBA API.	1740
BTAZC0295E An HBA API error occurred: no such binding.	1740
BTAZC0296E An HBA API error occurred: not a target.	1740
BTAZC0297E Unsupported FC4 with HBA API.	1740
BTAZC0298E Incapable error with the HBA API.	1741
BTAZC0299E An HBA API error occurred: target busy.	1741
BTAZC0301E An error occurred during a zoning command.	1741
BTAZC0302E An error occurred during zone control operation.	1741
BTAZC0303E CIM native error: Not Supported	1742
BTAZC0304E CIM native error: Unspecified Error	1742
BTAZC0305E CIM native error: Time Out	1742
BTAZC0306E CIM native error: Failed	1742
BTAZC0307E CIM native error: Invalid Parameter	1743
BTAZC0308E CIM native error: Access Denied	1743
BTAZC0309E CIM native error: Not Found	1743
BTAZC0310E CIM native error: Already Exist	1743

BTAZC0311E CIM native error: Insufficient Resources	1744
BTAZC0312E CIM native error: Empty Object Invalid	1744
BTAZC0313E CIM native error: No Transaction	1744
BTAZC0314E CIM native error: Transaction already on	1744
BTAZC0315E CIM native error: Cannot Commit Empty Objects	1745
BTAZC0316E CIM native error: Zone Database Full	1745
BTAZC0317E CIM native error: Too Many Members	1745
BTAZC0318E CIM native error: Fabric is busy	1745
BTAZC0319E Failed to create the zone zone_name on SAN SAN_name.	1746
BTAZC0320E Failed to create the zone alias zone_alias_name on SAN SAN_name.	1746
BTAZC0321E Failed to create the zone set zone_set_name on SAN SAN_name.	1746
BTAZC0322E Failed to delete the zone zone_name on SAN SAN_name.	1746
BTAZC0323E Failed to delete the zone alias zone_alias_name on SAN SAN_name.	1747
BTAZC0324E Failed to delete the zone set zone_set_name on SAN SAN_name.	1747
BTAZC0325E Failed to delete the zone member zone_member_name on SAN SAN_name.	1747
BTAZC0326E Failed to add the zone zone_name to zone set zone_set_name on SAN SAN_name.	1748
BTAZC0327E Failed to add the zone member zone_member_name to zone zone_name on SAN SAN_name.	1748
BTAZC0328E Failed to add the zone member zone_member_name to zone alias zone_alias_name on SAN SAN_name.	1748
BTAZC0329E Failed to add the zone alias zone_alias_name to zone zone_name on SAN SAN_name.	1748
BTAZC0330E Failed to remove the zone zone_name from zone set zone_set_name on SAN SAN_name.	1749
BTAZC0331E Failed to remove the zone member zone_member_name from zone zone_name on SAN SAN_name.	1749
BTAZC0332E Failed to remove the zone member zone_member_name from zone alias zone_alias_name on SAN SAN_name.	1749
BTAZC0333E Failed to remove the zone alias zone_alias_name from zone zone_name on SAN SAN_name.	1749
BTAZC0334E Failed to create the zone member zone_member_name on SAN SAN_name.	1750
BTAZC0335E Failed to activate Zone Set zone_set_name on SAN SAN_name.	1750
BTAZC0336E Failed to deactivate Zone Set zone_set_name on SAN SAN_name.	1750
BTAZC0337E Failed to enumerate the CIM entity AdminDomain for SAN SAN_name.	1751
BTAZC0338E Failed to start a Transaction for SAN SAN_name.	1751
BTAZC0339E Failed to commit a Transaction for SAN SAN_name.	1751
BTAZC0340E Failed to close the Session for SAN SAN_name.	1751
BTAZC0341E Failed to rollback a Transaction for SAN SAN_name.	1752
BTAZC0342E This command is not supported for the Fabric Agent.	1752
BTAZC0343E CIM native error: Transaction not available	1752
BTAZC0344E Zone must be included in ZoneSet. Zone Name.	1752
BTAZC0345E CIM error: Save ZoneDB To Switch Failed.	1753
BTAZC0346E CIM error: Save ZoneDBInfo Failed.	1753
BTAZC0347E CIM error: Zone Database Error.	1753
BTAZC0348E CIM error: Transaction Start Failed.	1753
BTAZC0349E CIM error: Transaction End Failed.	1754
BTAZC0350E CIM error: Transaction Terminate Failed.	1754
BTAZC0351E CIM error: Activate ZoneSet Failed.	1754
BTAZC0352E CIM error: Deactivate ZoneSet Failed.	1754
BTAZC0353E Unable to access the database to persist zoning changes.	1755
BTAZC5000I Started Zone Control layer.	1755
BTAZC5001I reserve: SAN=SAN_name, Agent=agent_name, Client=client_name, Token=token_ID.	1755
BTAZC5002I startTransaction: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1755
BTAZC5003I commitTransaction: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1755
BTAZC5004I rollbackTransaction: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1755
BTAZC5005I setZoneInfo: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1756
BTAZC5006I sendCommandArray: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1756
BTAZC5007I release: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1756
BTAZC5008I createZoneSet: zoneSetName=zone_set_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1756
BTAZC5009I createZone: zoneName=zone_name, zoneType=zone_type, zoneSubType=zone_subtype, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1756
BTAZC5010I createZoneAlias: zoneAliasName=zone_alias_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1756
BTAZC5011I createZoneMemberSettingData: zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, targetType=target_type, targetName=target_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1756
BTAZC5012I addZoneToZoneSet: zoneSetName=zone_set_name, zoneName=zone_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1757
BTAZC5013I addZoneMemberToZone: zoneName=zone_name, zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1757
BTAZC5014I addZoneAliasToZone: zoneName=zone_name, zoneAliasName=zone_alias_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1757
BTAZC5015I addZoneMemberToZoneAlias: zoneAliasName=zone_alias_name, zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1757
BTAZC5016I removeZoneFromZoneSet: zoneSetName=zone_set_name, zoneName=zone_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1757
BTAZC5017I removeZoneAliasFromZone: zoneName=zone_name, zoneAliasName=zone_alias_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1757
BTAZC5018I removeZoneMemberFromZone: zoneName=zone_name, zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1758
BTAZC5019I removeZoneMemberFromZoneAlias: zoneAliasName=zone_alias_name, zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.	1758

BTAZC5020I deleteZoneMember: zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name, Token=token_ID. result=return_code.	1758
BTAZC5021I deleteZone: zoneName=zone_name, SAN=SAN_name, Client=client_name, Token=token_ID. result=return_code.	1758
BTAZC5022I deleteZoneSet: zoneSetName=zone_set_name, SAN=SAN_name, Client=client_name, Token=token_ID. result=return_code.	1758
BTAZC5023I deleteZoneAlias: zoneAliasName=zone_alias_name, SAN=SAN_name, Client=client_name, Token=token_ID. result=return_code.	1758
BTAZC5024I activateZoneSet: zoneSetName=zone_set_name, SAN=SAN_name, Client=client_name, Token=token_ID. result=return_code.	1758
BTAZC5025I deactivateZoneSet: zoneSetName=zone_set_name, SAN=SAN_name, Client=client_name, Token=token_ID. result=return_code.	1759
BTAZC5026I ping: SAN=SAN_name, Agent=agent_name,	1759
BTAZC5027I readCurrentZoneDefinition: SAN=SAN_name, Client=client_name, Token=token_ID result=return_code.	1759
BTM - Common Information Model (CIM) agent messages	1759
BTM0001E Unable to connect to the SMI-S provider.	1763
BTM0002E CIM intrinsic method failure: value.	1764
BTM0003E Unable to disconnect from the SMI-S provider.	1764
BTM0004E Error getting Host Initiators connected to Target FCPort: value.	1764
BTM0005E Error getting Storage System's FCPorts: value.	1764
BTM0006E Error encountered while attempting SMI-S provider discovery.	1765
BTM0007E value is not a supported protocol for WBEM.	1765
BTM0008E Error getting storage systems from SMI-S provider at value, port value.	1765
BTM0009E Unable to get CIM_Product instance for this object: name.	1765
BTM0010E Unsupported Profile.	1765
BTM0011E Error getting Volumes for Storage System: storage system.	1766
BTM0012E Error getting paths from Hosts to Volumes for Storage System: name.	1766
BTM0013E Error getting detailed information for Storage System: value.	1766
BTM0014E Unable to create CIMObjectPath from String: name.	1766
BTM0015E Error getting Storage Pools for Storage System: value.	1767
BTM0016E Logical subsystems is an IBM-only concept.	1767
BTM0017E Error getting Storage System's logical subsystems: value.	1767
BTM0018E Error getting Storage System's Disk Groups: value.	1767
BTM0019E Error getting Storage Pools for this logical subsystem: name.	1768
BTM0020E Error getting Disks for this Storage Pool: value.	1768
BTM0021E Error getting Disks for this Disk Group: value.	1768
BTM0022E Error getting Volumes for this Storage Pool: value.	1768
BTM0023E Error enumerating namespaces.	1769
BTM0024E Error getting a specific Storage Pool: value.	1769
BTM0025E Unable to connect to SMI-S provider, bad/missing truststore or incorrect truststore password for SMI-S provider at value.	1769
BTM0026E Unable to connect to SMI-S provider, cannot find correct certificate in truststore for SMI-S provider at value.	1769
BTM0027E Unable to connect to SMI-S provider. Username, password, and/or protocol may be invalid for SMI-S provider at value.	1770
BTM0028E Unable to contact SMI-S provider at value. SMI-S provider may not be running.	1770
BTM0029E CIMService's hostname or IP is null.	1770
BTM0030E CIMService's port is invalid or null.	1770
BTM0031E CIMService's protocol is null.	1770
BTM0032E CIMAccessParameterSet's Certificate filename is invalid.	1771
BTM0033E CIMAccessParameterSet's password is null.	1771
BTM0034E CIMAccessParameterSet's username is null.	1771
BTM0035E string must be a CIM ObjectPath String for a Storage System.	1771
BTM0036E Invalid set of Volumes. Unable to get PathToLUNs.	1771
BTM0037E string must be a CIM ObjectPath String for a Storage Pool.	1772
BTM0038W Unable to determine Vendor of Storage System: name.	1772
BTM0039E Unable to determine the RAID Level of Volume: name.	1772
BTM0040W Unable to get CIM_Product info for Storage System: name.	1772
BTM0041E Unable to get CIM_Product info for Storage System. More than one Chassis associated to Storage System: name.	1772
BTM0042W Unable to get CIM_Product info for Storage System. No CIM_Product associated to Storage System's Chassis: value.	1773
BTM0043W More than one CIM_Product indirectly associated to Storage System: value.	1773
BTM0044E Unable to create Volume object: name.	1773
BTM0045E Unable to get Host Initiators that can access this volume: value.	1773
BTM0046E Host's permission value not recognized: value.	1774
BTM0047E Unable to get Disks for this Storage System: name.	1774
BTM0048E More than one Disk Group exists for this Disk: value.	1774
BTM0049E No Disk Group associated to this Disk: value.	1774
BTM0050E Unable to get Disk Group for Disk: value.	1775
BTM0051E This Volume is a component of more than one Storage System: value.	1775
BTM0052E This Volume is not part of a Storage System: name.	1775
BTM0053E Unable to get Volume: value.	1775
BTM0054E Please verify that you are running a supported version of a Common Information Model Agent for the storage subsystem.	1776
BTM0055E Unable to return systems associated with cluster: cluster.	1776
BTM0056E Unable to return Vendor of Cluster: cluster.	1776
BTM0057E Error getting Cluster backend controllers: controllers.	1776
BTM0058E cluster must be a CIM Object Path String for a Cluster.	1777
BTM0059E Backend Controllers not supported for vendor vendor on Cluster cluster.	1777
BTM0060E Unable to get Backend Controllers for this Cluster: cluster.	1777
BTM0061E Error getting Cluster managed disks: cluster.	1777
BTM0062E Unable to get Managed Disks for this Cluster: cluster.	1778

BTM0063E Unable to get Backend Controllers for Backend Volume volume on Cluster cluster.	1778
BTM0064E Unable to create Managed Disk object: disk	1778
BTM0065E Error getting Managed Disk Group Managed Disks: group	1778
BTM0066E Backend Volumes not supported for vendor vendor on Cluster cluster.	1779
BTM0067E Unable to get Managed Disks for this Managed Disk Group: group.	1779
BTM0068E No clusters associated with this Managed Disk Group: group.	1779
BTM0069E Error getting Virtual Disk managed disks: disk.	1779
BTM0070E volume must be a CIM Object Path String for a volume.	1780
BTM0071E Unable to get Managed Disks for this Virtual Disk: volume.	1780
BTM0072E No Clusters associated with this Virtual Disk: volume.	1780
BTM0073E Error getting Backend Controller managed disks: disk.	1780
BTM0074E controller must be a CIM Object Path String for a SCSI Controller.	1781
BTM0075E Unable to get Managed Disks for this Backend Controller: controller.	1781
BTM0076E No Clusters associated with this Backend Controller: controller.	1781
BTM0077E Error getting Cluster managed Disk Groups: cluster.	1781
BTM0078E Unable to get Managed Disk Groups for this Cluster: cluster.	1782
BTM0079E Unable to create Managed Disk Group Object: group.	1782
BTM0080E Error getting Managed Disk Group Virtual Disks: group.	1782
BTM0081E Unable to create Virtual Disk object: volume.	1782
BTM0082E Unable to get Cluster virtual disks: cluster.	1782
BTM0083E Unable to get Virtual Disks for this Cluster: cluster.	1783
BTM0084E Error getting Cluster: cluster.	1783
BTM0085E Error getting Storage System Type for Computer System: system.	1783
BTM0086E Error checking Storage system Level for Computer System: system.	1783
BTM0087E Unable to get the Storage System for this volume: volume.	1784
BTM0088E volume must be a CIM Object Path String for a Volume.	1784
BTM0089E Detected an unsupported level of the Common Information Model agent.	1784
BTM0090E Unable to create CIM Object Path String from Class Definition: class.	1784
BTM0091E Unable to determine the Privilege for Host Initiator value to access the Volume value.	1785
BTM0092W Cannot get Disk Drives for Storage Pool. No Storage Extents found for this Storage Pool: value.	1785
BTM0093E No Storage Extents found for this Disk Drive: value.	1785
BTM0094E This SMI-S provider version is not supported.	1785
BTM0095E This SMI-S provider vendor is not supported.	1786
BTM0096E Unable to retrieve LSI SMI-S CIM provider version.	1786
BTM0098E Unable to retrieve CIM Object Path for Storage System: storage system from the SMI-S provider.	1786
BTM0100E Cannot find unassigned LUNs because the storage pool list is null.	1787
BTM0101E Unable to retrieve the Storage System path.	1787
BTM0102E Unable to retrieve the Hardware Account path.	1787
BTM0103E The Storage System path is null.	1787
BTM0104E The Hardware Account path is null.	1787
BTM0105E Unable to retrieve the FC Port path.	1788
BTM0106E The FC Port path is null.	1788
BTM0107E Unable to retrieve the Authorization Service path for Subsystem: value.	1788
BTM0108E The Authorization Service path is null for Subsystem: value.	1788
BTM0109E There are multiple Authorization Service paths for the Subsystem: value.	1789
BTM0110E Unable to retrieve the Main Controller path for Subsystem: value.	1789
BTM0111E The Main Controller path is null for Subsystem: value.	1789
BTM0112E There are multiple Main Controller paths for the Subsystem: value.	1789
BTM0113E Unable to retrieve Clone Controller path for the Subsystem: value Hardware Account: value FC Port: value.	1790
BTM0114E The Clone Controller path is null for Subsystem: value Hardware Account: value FC Port: value.	1790
BTM0115E Unable to retrieve the Hardware Account for the Clone Controller: value.	1790
BTM0116E The Hardware Account for the Clone Controller: value is null.	1790
BTM0117E No Hardware Account for the Clone Controller: value.	1791
BTM0118E No Access Control Information for the Clone Controller: name.	1791
BTM0119E Unable to retrieve the FC Port for the Clone Controller: name.	1791
BTM0120E The FC Port for the Clone Controller: name is null.	1791
BTM0121E Unable to create a Clone Controller with FC Port: port number Authorization Service: service.	1792
BTM0122E Unable to remove Clone Controller: controller name.	1792
BTM0123E Unable to Assign Access with Hardware Account: account number Clone Controller: controller Authorization Service: service.	1792
BTM0124E Unable to Remove Access with Hardware Account: account number Clone Controller: controller Authorization Service: service.	1792
BTM0125E Unable to Attach Volume with Volume volume name Clone Controller: controller.	1793
BTM0126E Unable to Detach Volume with Volume volume name Clone Controller: controller.	1793
BTM0127E Unable to get Volume, Subsystem, or AuthorizationService path.	1793
BTM0128E Unable to assign Volume value to Path [name, name] on Subsystem name using Controller name with Authorization Service name.	1793
BTM0129E Unable to unassign Volume name to Path [name, name] on Subsystem name using Controller name with Authorization Service name.	1794
BTM0130E Rolling back value assignments.	1794
BTM0131E Rolling back value unassignments.	1794
BTM0132E Error getting unassigned LUNs.	1794
BTM0133E Error assigning paths.	1795
BTM0134E Error unassigning paths.	1795
BTM0141E Unable to Attach Volume with Volume Storage Volume for Controller Controller using Controller Configuration Service: Controller Configuration Service.	1795

BTM0142E Unable to Detach Volume with Volume Storage Volume for Controller Controller using Controller Configuration Service: Controller Configuration Service.	1795
BTM0149E Error calling extrinsic method {0} rc = {1}: Invalid Storage Pool There are multiple Privilege Management Service paths for the Subsystem: Storage Subsystem.	1796
BTM0151E There are multiple Privilege Management Service paths for the Subsystem: Storage Subsystem.	1796
BTM0152E The Privilege Management Service path is null for Subsystem: Storage Subsystem.	1796
BTM0153E There are multiple Controller Configuration Service paths for the Subsystem: Storage Subsystem.	1797
BTM0154E The Controller Configuration Service path is null for Subsystem: Storage Subsystem.	1797
BTM0155E Unable to assign Volume Storage Volume to Path [Hardware Account, FC Port] on Subsystem Storage Subsystem using Controller Controller with Privilege Management Service Privilege Service and Controller Configuration Service Controller Service.	1797
BTM0156E Unable to unassign Volume Storage Volume to Path [Hardware Account, FC Port] on Subsystem Storage Subsystem using Controller Controller with Privilege Management Service Privilege Service and Controller Configuration Service Controller Service.	1797
BTM0157E Unable to retrieve the model volume path.	1798
BTM0158E Unable to assign volume for an invalid client request.	1798
BTM0159E Unable to unassign volume for an invalid client request.	1798
BTM0200E Unable to create Storage Volume of size value in Storage Pool value.	1798
BTM0201E Storage Volume of size value not created in Storage Pool value.	1799
BTM0202E Unable to retrieve Storage Service for Storage Pool value.	1799
BTM0203E Unable to retrieve Storage System for Storage Pool value.	1799
BTM0204E Storage Pool used to create the Storage Volume of size value is null.	1799
BTM0205E Size used to create the Storage Volume on Storage Pool value is null.	1800
BTM0206E Both the Storage Pool and the size to create the Storage Volume are null.	1800
BTM0207E Storage Volume identification is null and Storage Volume cannot be located.	1800
BTM0208E Storage Volume identification value failed to retrieve Storage Volume.	1800
BTM0209E Storage Volume identification value cannot be used to locate a Storage Volume.	1801
BTM0210E Storage Volume object is null for Storage System value.	1801
BTM0211E Storage System is null for Storage Volume value.	1801
BTM0212E Both the Storage System and the Storage Volume object are null.	1801
BTM0213E Unable to return the Paths to Storage Volume value on Storage System value.	1802
BTM0214E There are no Paths to Storage Volume value on Storage System value.	1802
BTM0215E Client connection is null when retrieving Storage Volume identification value.	1802
BTM0216E Storage Volume identification is null.	1802
BTM0217E Both the Client connection and the Storage Volume identification are null.	1803
BTM0218E Unable to retrieve Storage Volume object using Storage Volume identification value.	1803
BTM0219E Storage System Type of value is not valid for Storage Volume identification value.	1803
BTM0220E Unable to locate Storage Volume object using Storage Volume identification value.	1803
BTM0221E Instance of Storage Volume is null.	1804
BTM0222E Unable to retrieve Storage Volume identification from Storage Volume instance.	1804
BTM0223E Retrieved invalid Storage System name of value from Storage Volume instance.	1804
BTM0224E List of Storage Volume objects is invalid.	1804
BTM0225E Unable to complete list of Storage Pool objects for Storage Volume value.	1805
BTM0226E Unable to complete list of Storage Pool objects without a Storage Volume object.	1805
BTM0227E No Storage Pool objects returned for Storage Volume value.	1805
BTM0228E Unable to enumerate Storage Pool objects for Storage Volume value.	1805
BTM0229E Unable to return Storage Pool objects for Storage Volume value.	1806
BTM0230E Unable to generate a list of Storage Pool objects for Storage Volume value.	1806
BTM0231E Unable to generate a list of Storage Pool objects without a Storage Volume object.	1806
BTM0232E Unable to create Storage Volumes.	1806
BTM0233E Unable to select Storage Pools.	1807
BTM0234E More than one Storage Service found for Storage System value.	1807
BTM0235E Failed to retrieve newly created Storage Volume of size value in Storage Pool value.	1807
BTM0236E Storage Volume to be removed is null.	1807
BTM0237E Storage Volume value is not removed.	1808
BTM0238E Failed to remove Storage Volume value.	1808
BTM0239E Unable to retrieve Storage Service for Storage Volume value.	1808
BTM0400E Error calling extrinsic method value rc = value: Unsupported method rc.	1808
BTM0401E Error calling extrinsic method value rc = value: Unknown error.	1809
BTM0402E Error calling extrinsic method value rc = value: Not Supported.	1809
BTM0403E Error calling extrinsic method value rc = value: Failed.	1809
BTM0404E Error calling extrinsic method value rc = value: Invalid parameter ports.	1809
BTM0405E Error calling extrinsic method value rc = value: Invalid controller.	1810
BTM0406E Error calling extrinsic method value rc = value: Missing required property within Subject or Target.	1810
BTM0407E Error calling extrinsic method value rc = value: Invalid parameter.	1810
BTM0408E Error calling extrinsic method value rc = value: Input controller must have AuthorizationView set to FALSE.	1810
BTM0409E Error calling extrinsic method value rc = value: Invalid LogicalDevice instance.	1811
BTM0410E Error calling extrinsic method value rc = value: Hardware implementation requires null DeviceNumber.	1811
BTM0411E Error calling extrinsic method value rc = value: Input size is bigger than the free spaces left in the InPool.	1811
BTM0412E Error calling extrinsic method value rc = value: Authorization failure.	1811
BTM0413E Error calling extrinsic method value rc = value: Cannot remove device because it is not attached.	1812
BTM0414E Error calling extrinsic method value rc = value: Invalid parameter Subject.	1812
BTM0415E Error calling extrinsic method value rc = value: Invalid StorageSetting.	1812
BTM0416E Error calling extrinsic method value rc = value: Invalid parameter Target.	1812

BTM0417E Error calling extrinsic method value rc = value: Input size is invalid, either less than or equal to 0, or is null.	1813
BTM0418E Error calling extrinsic method value rc = value: Access is not yet assigned.	1813
BTM0419E Error calling extrinsic method {0} rc = {1}: Invalid Storage Pool method rc = return code: Invalid Storage Pool.	1813
BTM0420E Error calling extrinsic method value rc = value: The specified Subject and Target are not associated.	1813
BTM0421E Error calling extrinsic method value rc = value: Should remove access first.	1814
BTM0422E Error calling extrinsic method value rc = value: Should assign access first.	1814
BTM0423E Error calling extrinsic method value rc = value: Element type should be 2, meaning Storage Volume. (The Volume input should be null.).	1814
BTM0424E Error calling extrinsic method value rc = value: Supports single target only.	1815
BTM0425E Error calling extrinsic method value rc = value: UserIDType should be PortWWN.	1815
BTM0426E Error calling extrinsic method value rc = value: Cannot create a temporary controller.	1815
BTM0427E Error calling extrinsic method value rc = value: LogicalDevice instance is already attached to a Host.	1815
BTM0428E Error calling extrinsic method value rc = value: Should detach the device first.	1816
BTM0429E Error calling extrinsic method value rc = value: UserIDType should be the same as Name, which is PortWWN.	1816
BTM0430E Error calling extrinsic method value rc = value: IBMTSESS cannot create view as specified.	1816
BTM0431E Error calling extrinsic method value rc = value: Controller processing fails (Failed to delete temporary controller from repository server.	1816
BTM0432E Error calling extrinsic method value rc = value: IBMTSESS cannot attach the device as specified (ESSCLI cannot create VolumeAccess as specified).	1817
BTM0433E Error calling extrinsic method value rc = value: IBMTSESS cannot remove the device as specified (ESSCLI cannot delete VolumeAccess as specified).	1817
BTM0434E Error calling extrinsic method value rc = value: IBMTSESS does not support modification of volume.	1817
BTM0435E Error calling extrinsic method value rc = value: IBMTSESS cannot AssignAccess as specified (ESSCLI cannot create or set HostConnection).	1817
BTM0436E Error calling extrinsic method value rc = value: IBMTSESS cannot RemoveAccess as specified (ESSCLI cannot delete or set HostConnection).	1818
BTM0437E Error calling extrinsic method value rc = value: HardwareAccount instance already exists or HardwareAccount processing fails.	1818
BTM0438E Error calling extrinsic method value rc = value: HardwareAccount processing fails.	1818
BTM0439E Error calling extrinsic method value rc = value: IBMTSESS cannot create volume as specified (ESSCLI cannot create volume).	1819
BTM0440E Error calling extrinsic method value rc = value: Controller processing failed.	1819
BTM0441E Error calling extrinsic method value rc = value: HardwareAccount processing failed.	1819
BTM0442E Error calling extrinsic method name rc = value: Creating indication failure.	1819
BTM0443E Error calling extrinsic method name rc = value: The requested logical subsystem already contains the maximum number of volumes allowed.	1820
BTM0444E Error calling extrinsic method value rc = value: The requested amount of volume addresses exceeds the maximum number of volumes allowed in the given logical subsystems.	1820
BTM0459E Unable to get CIM_StorageExtent instance for this object: disk drive from the SMI-S provider.	1820
BTM0460E Unable to get CIM_PhysicalPackage instance for this object: physical package from the SMI-S provider.	1820
BTM0461E Unable to get CIM_SoftwareIdentity instance for this object: software identity from the SMI-S provider.	1821
BTM0462E Error calling extrinsic method method rc = return code: Invalid Protocol.	1821
BTM0463E Error calling extrinsic method method rc = return code: Cannot create temporary controller in SMI-S provider repository.	1821
BTM0464E Unable to retrieve CIM_SystemSpecificCollection paths for CIM_Privilege: CIM Object.	1821
BTM0465E Unable to retrieve CIM_StorageHardwareID paths for CIM_SystemSpecificCollection: CIM Object.	1822
BTM0466E Unable to retrieve CIM_StorageHardwareID paths for CIM_Privilege: CIM Object.	1822
BTM0467E Unable to retrieve CIM_Privilege paths for CIM_SCSIProtocolController: CIM Object.	1822
BTM0468E Unable to retrieve CIM_SCSIProtocolController paths for CIM_StorageVolume: CIM Object.	1822
BTM0469E Cannot determine if CIM_StorageVolume has been surfaced: CIM Object.	1823
BTM0470E Cannot retrieve cache size for CIM_ComputerSystem: CIM Object.	1823
BTM0550W Cannot get Disk Drives for Storage Pool. No Disk Drives found for this Storage Pool: value.	1823
BTM0551W Cannot get Disk Drives for this Storage Pool: value.	1823
BTM0552W Value of value not available for: property .	1824
BTM0553I Probing Disks for DiskGroup: value	1824
BTM0554I Probing Disks for StoragePool: value	1824
BTM0555I Number of Disks Found Currently: value. {0}. Continuing to Probe Disks.	1824
BTM0556W Cannot get Disk Drives for Storage System. No Disk Drives found for this Storage System: value.	1824
BTM0557W Cannot get Disk Drives for this Storage System: value.	1824
BTM0558I Number of Volumes Found Currently: value. Continuing to Probe Volumes.	1825
BTM0559I Probing Volumes for StoragePool: value.	1825
BTM0560I Probing Volumes for Storage System: value.	1825
BTM0561I Probing Disks for Storage System: value.	1825
BTM0562I Probing Storage Pools for Storage System: value.	1825
BTM0563I Probing properties of Storage System: value.	1825
BTM0564W More than one CIM_Product indirectly associated to the following Device: value.	1825
BTM0565W Exception caught while getting CIM_Product info for Storage System: value.	1826
BTM0566W Unable to get CIM_Product info for Device: value.	1826
BTM0567W Exception caught while getting CIM_Product info for Device: value.	1826
BTM0568I value Volumes Found.	1826
BTM0569I value Disks Found.	1826
BTM0571W Exception caught while getting Host Initiators that can access this volume: value.	1827
BTM0572W Exception caught while trying to determine RAID Level for StoragePool: value.	1827
BTM0573E Exception caught while formatting this Host Bus Adapter port World Wide Name: value.	1827
BTM0574W Capacity of Disk Drive is not available: value.	1827
BTM0575W Exception caught while getting Host Initiators access to Volumes through this View: value.	1828
BTM0576I Probing Views of Host Initiator access to Volumes.	1828
BTM0577I value Views Found.	1828
BTM0578E Unable to connect to SMI-S provider. None of the default namespaces are valid for this SMI-S provider.	1828
BTM0600E Unable to get Array Site for Disk: value.	1828
BTM0601E Error calling extrinsic method value rc = value: A timeout occurred trying to call the method.	1829
BTM0602E Error calling extrinsic method value rc = value: The instance of the Logical Device is invalid.	1829

BTM0603E Error calling extrinsic method value rc = value: There is a conflict in the Device Number.	1829
BTM0604E Error calling extrinsic method value rc = value: A Device Number parameter must be provided.	1829
BTM0605E Error calling extrinsic method value rc = value: A null Device Number is required by the device.	1830
BTM0606E Error calling extrinsic method value rc = value: The device is busy.	1830
BTM0607E Error calling extrinsic method value rc = value: The Protocol Controller is invalid.	1830
BTM0608E Error calling extrinsic method value rc = value: The volume types are invalid.	1830
BTM0609E Error calling extrinsic method value rc = value: One or more parameters are in the wrong System Scope.	1831
BTM0610E Error calling extrinsic method value rc = value: The controller needs to be created first.	1831
BTM0611E Error calling extrinsic method value rc = value: The ESSCLI call to create the volume access failed.	1831
BTM0612E Error calling extrinsic method value rc = value: The ESSCLI call to list the volume access failed.	1831
BTM0613E Error calling extrinsic method value rc = value: The instance of the Logical Device is not associated with the Controller.	1832
BTM0614E Error calling extrinsic method value rc = value: The subject is not supported.	1832
BTM0615E Error calling extrinsic method value rc = value: The Privilege is not supported.	1832
BTM0616E Error calling extrinsic method value rc = value: The Target is not supported.	1832
BTM0617E Error calling extrinsic method value rc = value: A null parameter is not supported.	1833
BTM0618E Error calling extrinsic method value rc = value: Configuration Service is in use.	1833
BTM0619E Error calling extrinsic method value rc = value: The size is invalid.	1833
BTM0620E Error calling extrinsic method value rc = value: The Element Type is invalid.	1833
BTM0621E Error calling extrinsic method value rc = value: The Goal is invalid.	1834
BTM0622E Error calling extrinsic method value rc = value: The Storage Pool is invalid.	1834
BTM0623E Error calling extrinsic method value rc = value: The redundancy for the Storage Pool is invalid.	1834
BTM0624E Error calling extrinsic method value rc = value: The requested Data Type does not match the Data Type for the Storage Pool.	1834
BTM0625E Error calling extrinsic method value rc = value: The Data Type is invalid.	1835
BTM0626E Error calling extrinsic method value rc = value: The Element is invalid.	1835
BTM0627E Error calling extrinsic method value rc = value: No parameters were specified for the modification.	1835
BTM0628E Error calling extrinsic method value rc = value: Unable to create volume.	1835
BTM0629E Error calling extrinsic method value rc = value: The LSS already contains the maximum number of volumes.	1836
BTM0630E Error calling extrinsic method value rc = value: There are not enough volume addresses in the LSS.	1836
BTM0631E Error calling extrinsic method value rc = value: The Identification parameter is missing or not unique.	1836
BTM0632E Error calling extrinsic method value rc = value: A null Ports parameter is required by the Controller.	1837
BTM0633E Error calling extrinsic method value rc = value: The Controller is busy.	1837
BTM0634E Error calling extrinsic method value rc = value: The Identity is invalid.	1837
BTM0635E Error calling extrinsic method value rc = value: The Element Name is invalid.	1837
BTM0636E Error calling extrinsic method value rc = value: The Protocol is invalid.	1838
BTM0637E Error calling extrinsic method value rc = value: The Privilege is invalid.	1838
BTM0638E Error calling extrinsic method value rc = value: The Ports are invalid.	1838
BTM0639E Error calling extrinsic method value rc = value: The host connection could not be deleted.	1838
BTM0640E Error calling extrinsic method value rc = value: The host connection could not be created.	1839
BTM0641E Error calling extrinsic method value rc = value: The host connection could not be set.	1839
BTM0642E Error calling extrinsic method value rc = value: No Ports are available in this configuration.	1839
BTM0701I Probing Managed Disks for Managed Disk Group: value	1839
BTM0702I Number of Managed Disks currently found: value. Continuing to probe managed disks.	1840
BTM0703I value Managed Disks found.	1840
BTM0704I Probing Virtual Disks for Cluster: value	1840
BTM0705I Number of Virtual Disks currently found: value. Continuing to probe Virtual Disks.	1840
BTM0706I value Virtual Disks found.	1840
BTM0707I Probing Virtual Disks for Managed Disk Group: value	1840
BTM0708I Probing Managed Disks for Cluster: value	1840
BTM0709I Probing Managed Disks for Backend Controller: value	1840
BTM0710E Unable to retrieve data for Managed Disk: value	1840
BTM0711E Unable to retrieve data for Managed Disk Group: value	1841
BTM0712E Unable to retrieve data for Virtual Disk: value	1841
BTM0713E Unable to retrieve data for Backend Controller: value	1841
BTM0714E Unable to retrieve data for FC Port: value	1841
BTM0715E Unable to retrieve data for value Managed Disk(s) among the value Managed Disks found.	1842
BTM0716E Unable to retrieve data for value Virtual Disk(s) among the value Virtual Disks found.	1842
BTM0717E Unable to retrieve FC Ports for Cluster: value	1842
BWN - Disk User Interface messages	1842
BWN000000E An object must be selected.	1843
BWN000200E The minimum size of the volume cannot be less than minimum volume size.	1843
BWN000201E The volume volume name cannot be deleted because there are host ports assigned to it.	1843
BWN000202E The volume volume name cannot be deleted, because volume deletion is not supported by this storage subsystem.	1843
BWN000203E No storage pool is available.	1844
BWN000204E The maximum size of the volume cannot be greater than maximum volume size.	1844
BWN000205E The maximum size of the XIV volume cannot be greater than maximum volume size. The size of the volume must be changed.	1844
BWN000206E The number of volumes must be reduced so that volume(s) with selected size can be created.	1844
BWN000207E The volume deletion failed. Please check job log for detailed information about the error.	1844
BWN000208E Invalid characters in volume name. Volume name can only contain: A-Z, a-z, 0-9, _, -, . and space.	1845
BWN000300E A valid quantity must be selected. It must be between 1 and max quantity.	1845
BWN000301E A valid size must be selected. It must be between 0 and max size.	1845
BWN000302E The maximum virtual disk size must be less than or equal to the available capacity (available capacity) for one virtual disk in the managed-disk group. If multiple virtual disks are to be created, the virtual-disk size must be less than or equal to the available capacity divided by the number of virtual disks	1845

BWN000303E The number of virtual disks is invalid. The maximum number of virtual disks that can be created is (available quantity).	1845
BWN000304E At least one managed disk must be selected.	1846
BWN000305E The virtual disk vdisk name cannot be deleted because there are host ports assigned to it.	1846
BWN000306E The virtual disk name is not valid.	1846
BWN000307E Sequential virtual disks and multiple managed disks are selected, but round-robin assignment is not specified.	1846
BWN000308E The number of selected mdisks must be equal to the number of vdisks to be created - vdisk no..	1846
BWN000309E The selected managed disks could not be added to the managed-disk group.	1847
BWN000310E The selected host type does not match the host type of the selected host ports to be assigned.	1847
BWN000311E The selected host ports have identical WWPNs. Select the host ports with different WWPNs.	1847
BWN000312E The maximum virtual-disk real size must be less than or equal to the available capacity (available capacity) in the managed-disk group. If multiple virtual disks are to be created, the virtual-disk real size must be less than or equal to the available capacity divided by the number of virtual disks.	1847
BWN000313E When creating Space Efficient virtual-disks, the maximum virtual-disk size must not exceed maximum size.	1848
BWN000314E The virtual-disk warning size must be greater than 0 and cannot exceed 100 percent.	1848
BWN000315W This SMI-S provider is already defined with the same parameters. Would you like to save it anyway?	1848
BWN000316W This IBM Spectrum Control Server is already defined with the same parameters. Would you like to save it anyway?	1848
BWN000317W Testing SMI-S provider connectivity can take up to several minutes in case of an incorrectly entered port number, network problems or an unpassed firewall. Would you like to continue anyway?	1848
BWN000318E No managed disk is found for the selected mdisk group.	1849
BWN000319E The length of the generated virtual disk name (virtual disk) exceeds the maximum permitted length (maximum length of characters).	1849
BWN000600E Some of the selected host ports do not have a host connection configured on the subsystem\and require a host type to be specified. Select the appropriate host type to be used.	1849
CMMNP - Command Line Interface (CLI) infrastructure messages	1849
CMMNP2001I Nothing to modify.	1850
CMMNP2002I Unsupported VALUE_0 command completed sucessfully.	1850
CMMNP2900I Command "VALUE_0" aborted.	1850
CMMNP4500W No VALUE_0 instances found in the system.	1850
CMMNP9002E Cannot modify. VALUE_0 "VALUE_1" does not exist.	1851
CMMNP9003E No VALUE_0 instancesVALUE_1 found that match criteria: VALUE_2.	1851
CMMNP9004E VALUE_0 "VALUE_1" does not exist.	1851
CMMNP9005E Unsupported VALUE_0 command failed with a value VALUE_1	1851
CMMUI4444E User name not specified.	1851
CMMUI9000E [3]An unknown value "VALUE_0" for command "VALUE_1" was returned.	1852
CMMUI9001E unknown	1852
CMMUI9006E [3]Command failed to execute correctly.	1852
CMMUI9007E [3]Password file access error: VALUE_0.	1852
CMMUI9008E [3]Malformed password file. First line of the file requires a colon delimited user:password string	1852
CMMUI9010E [1]Invalid command: "VALUE_0" not found.	1852
CMMUI9011E [21]Invalid flag: "VALUE_0".	1853
CMMUI9012E [21]Value "VALUE_0" for flag "-VALUE_1" is formatted incorrectly.	1853
CMMUI9013E [21]Missing parameter specifier after "-"	1853
CMMUI9014E [21]Flag "VALUE_0" already specified.	1853
CMMUI9015E [21]Flag "VALUE_0" missing required value.	1853
CMMUI9016E [21]Invalid value for VALUE_0: VALUE_1.	1854
CMMUI9017E [21]The VALUE_0 flag cannot be used when the VALUE_1 option is specified.	1854
CMMUI9018E [21]Command "VALUE_0" formatted incorrectly	1854
CMMUI9019E [21]Missing required parameter: "VALUE_0"	1854
CMMUI9020E [21]"VALUE_0" is mutually exclusive of "VALUE_1"	1854
CMMUI9021E [21]VALUE_0 exceeds the maximum allowable value of VALUE_1 for parameter "VALUE_2"	1854
CMMUI9022E [21]VALUE_0 does not meet the minimum allowable value of VALUE_1 for parameter "VALUE_2"	1855
CMMUI9023E [21]Unmatched VALUE_0 characters	1855
CMMUI9024E [21]Invalid value for VALUE_0: exceeds VALUE_1 characters	1855
CMMUI9025E [21]Value "VALUE_0" for argument "VALUE_1" invalid	1855
CMMUI9026E VALUE_0 "VALUE_1" does not exist.	1855
CMMUI9027E [21]Value "VALUE_0" cannot be accepted with any other value for the "-VALUE_1" flag.	1856
CMMUI9028E [3]The help page for command "VALUE_0" does not exist.	1856
CMMUI9029E [21]It is required to specify parameter "VALUE_1" when using parameter "VALUE_0"	1856
CMMUI9030E File "VALUE_0" doesn't exist.	1856
CMMUI9031E [21]Parameter "VALUE_0" cannot be used in the same command as parameter "VALUE_1".	1856
CMMUI9032E VALUE_0 "VALUE_1" already exists.	1856
CMMUI9033E [21]Value "VALUE_0" for flag "-VALUE_1" already specified.	1857
CMMUI9034E [21]Multiple targets not allowed for command "VALUE_0"	1857
CMMUI9035E [21]You cannot specify multiple VALUE_0s when using the VALUE_1 flag.	1857
CMMUI9036E [21]Invalid value "VALUE_1" for "VALUE_0": contains unsupported characters.	1857
CMMUI9037E [21]Invalid "VALUE_0" name "VALUE_1": contains unsupported characters.	1857
CMMUI9038E [21]Invalid value for VALUE_0: value other than VALUE_1 or VALUE_2 specified.	1858
CMMUI9039E [21]Value for flag "-VALUE_0" can not contain a "VALUE_1".	1858
CMMUI9040E [21]Number of entries (VALUE_0) is exceeded for the "-VALUE_1" flag.	1858
CMMUI9041E [21]Entry "VALUE_0" exceeds the length limit (VALUE_1) for one item for the "-VALUE_2" flag.	1858
CMMUI9042E [21]Value for -VALUE_0 must be VALUE_1 the current setting of VALUE_2.	1858
CMMUI9043E [21]Unrecognized syntax error in command "VALUE_0"	1859
CMMUI9044E Cannot run "VALUE_0" as a command within the VALUE_1 application.Tip: Enter "help VALUE_2" for more information.	1859
CMMUI - CIM Object Manager messages	1859

CMMUI9900E User access to CIMOM server denied.	1859
CMMUI9901E User access to command "VALUE_0" denied.	1859
CMMUI9902E Invalid key in truststore.	1860
CMMUI9903E The IBM CIM Object Manager has encountered an internal error.	1860
CMMUI9904E Truststore access failure.	1860
CMMUI9905E Namespace not found in the CIMOM server: "VALUE_0".	1860
CMMUI9906E Host url unspecified to CIMOM server.	1860
CMMUI9907E Invalid host specified to CIMOM server: "VALUE_0".	1860
CMMUI9908E Could not connect to CIMOM server.	1861
CMMUI9909E Invalid port address for CIMOM server: "VALUE_0".	1861
CMMUI9910E An unexpected CIMOM based error occurred: "VALUE_0".	1861
CNFG - Spectrum Control Configuration messages	1861
CNFG00001E The Prefix can not be blank, or contain any of the following characters: \\\\:.*?>< ."	1862
CNFG00002E The prefix {0} can not be used because {1} duplicate names have been found. For example: {2}	1862
CNFG00003E Invalid HMC Address.	1862
CNFG00004E Invalid Username for subsystem {0}.	1862
CNFG00005E Invalid Password for subsystem {0}.	1862
CNFG00006E Invalid IP Address for subsystem {0}.	1862
CNFG00007E Invalid Admin Username for subsystem {0}.	1863
CNFG00008E Invalid Admin Password for subsystem {0}.	1863
CNFG00009E Invalid Username for subsystem {0}.	1863
CNFG00010E Invalid Public SSH Key for subsystem {0}.	1863
CNFG00011E Invalid Private SSH Key for subsystem {0}.	1863
CNFG00012E Invalid SMI-S provider Host.	1864
CNFG00013E Invalid SMI-S provider Namespace.	1864
CNFG00014E Invalid SMI-S provider Port Number.	1864
CNFG00015E Invalid Out of Band Agent Host name.	1864
CNFG00016W You will need to update the credentials for {0}. Do you want to update the credentials now?	1864
CNFG00017E Connection test failed with status: {0}	1864
CNFG00018E Connection test to SMI-S provider {0} failed with status: {1}. {2}	1865
CNFG00019E The server was unable to contact out of band fabric agent at address {0} for the following reason: {1}	1865
CNFG00020E The filter text [{0}] contains invalid characters. Click the help button for allowable characters	1865
CNFG00021E Only {0} log files may be opened at any one time. Please reduce the number of selected log files and try again.	1865
CNFG00022E Unable to retrieve job data due to an internal error. Please check Data server logs.	1865
CNFG00023E Invalid Public SSH Key for Storwize V7000U File Module {0}.	1866
CNFG00024E Invalid Username for Storwize V7000U File Module {0}.	1866
CNFG00025E Invalid Password for Storwize V7000U File Module {0}.	1866
CNFG00026E Invalid Passphrase for Storwize V7000U File Module {0}.	1866
CNFG00027E Connection failed for Storwize V7000U File Module. Check Management Console {0}.	1866
CNFG00028E Invalid Public SSH Key for IBM IBM SONAS Device {0}.	1867
CNFG00029E Invalid Username for IBM SONAS Device {0}.	1867
CNFG00030E Invalid Password for IBM SONAS Device {0}.	1867
CNFG00031E Invalid Passphrase for IBM SONAS Device {0}.	1867
DIS - Discovery messages	1867
DIS0001E Discover: invalid command " command".	1867
DIS0001I Command command selected.	1868
DIS0002E Discover: invalid option " option".	1868
DIS0003E Discover: value missing for option " option".	1868
DIS0004E Command command is missing a required parameter.	1868
DIS0005E NetWare functionality implemented only on Windows	1869
DIS0006E Error occurred while processing datafile datafile.	1869
DIS0007E Unable to send Discovery results to the server.	1869
DIS0008E NAS Server server name was not found as a referenced computer in the mnttab of computer computer.	1869
EMSG - DS8000 management console messages	1870
EMSG0001E The DS8000 Element Managers view is not accessible because the Device Server is down.	1870
EMSG0002E The embedded browser widget was unexpectedly destroyed. Click OK to reset display.	1870
EMSG0003I The element manager is about to be removed. Once removed the element manager will not be accessible from IBM Spectrum Control. Do you wish to continue?	1871
EMSG0004I Connection test to the element manager element manager passed.	1871
EMSG0005E Connection test to the element manager element manager failed.	1871
EMSG0006I SMI-S provider connection was removed from element manager.	1871
EMSG0007E A problem occurred removing the SMI-S provider from the element manager.	1871
EMSG0008I The SMI-S provider has been added successfully. IBM Spectrum Control has started the discovery job for the device managed by this SMI-S provider. To check the status of the jobs, go to the IBM Spectrum Control perspective and check the following navigation tree nodes: Administrative Services -> Discovery -> CIMOM	1872
EMSG0009E The element manager already exists.	1872
EMSG0010E A problem occurred adding the element manager to IBM Spectrum Control.	1872
EMSG0011E Connection test to Element Manager failed.	1872
EMSG0012E Unable to establish an https connection to the element manager.	1872
EMSG0013E SMI-S provider connection was added to the element manager, however the SMI-S provider discovery job failed to launch.	1873
EMSG0014E An element manager named manager.name already exists.	1873
EMSG0015E A problem occurred updating the element manager to IBM Spectrum Control.	1873

EMSG0016E Connection test to SMI-S provider FAILED due to status code.	1873
EMSG0017E The element manager's URL is not in the correct format.	1874
EMSG0018E A problem occurred locating the element manager in IBM Spectrum Control.	1874
EMSG0019E A problem occurred locating the SMI-S provider associated with the element manager.	1874
EMSG0020E Unable to reset DS8000 Element Manager password.	1874
EMSG0021E To be added to IBM Spectrum Control, all DS8000 Element Manager software is prior to release 4.2 need Username and password.	1875
EMSG0022E User credentials provided for DS8000 Element Manager are not valid. Provide valid credentials.	1875
EMSG0023E Unable to Add or Modify DS8000 Element manager for unknown reasons.	1875
EMSG0024E Unable to Add or Modify DS8000 Element manager. Provide right password.	1875
EMSG0025E Unable to Add or Modify DS8000 Element manager. Account does not exist.	1876
EMSG0026E Unable to Add or Modify DS8000 Element manager. User Account is locked.	1876
EMSG0027E Unable to Add/Modify DS8000 Element manager. Storage Authentication Service (SAS) database cannot be accessed.	1876
EMSG0028E Unable to Add or Modify DS8000 Element manager. Parameters passed are not valid.	1876
EMSG0029E Unable to Add or Modify DS8000 Element manager. Storage Authentication Service(SAS) database login task has failed. Username or password is Incorrect.	1877
EMSG0030E Unable to Add or Modify DS8000 Element manager. The external user account or user group is not mapped to a DS series user role.	1877
EMSG0031E Unable to Add or Modify DS8000 Element manager. Token submitted to Authentication policy is not supported.	1877
EMSG0032E Unable to Add or Modify DS8000 Element manager. Token submitted for authentication has expired. Re-authenticate to continue.	1877
EMSG0033E Unable to Add or Modify DS8000 Element manager. No URL provided for Storage Authentication Service (SAS) policy.	1878
EMSG0034E Unable to Add or Modify DS8000 Element manager. The host specified in the URL is not a known host or is not reachable.	1878
EMSG0035E Unable to Add or Modify DS8000 Element manager. The specified truststore does not have a valid certificate for the Storage Authentication Service (SAS).	1878
EMSG0036E Unable to Add or Modify DS8000 Element manager. Connection request to the Storage Authentication Service (SAS) is refused.	1878
EMSG0037E Unable to Add or Modify DS8000 Element manager. Connection request to the Storage Authentication Service (SAS) has failed due to socket timeout.	1879
EMSG0038E Unable to Add or Modify DS8000 Element manager. Failed to login to DS8000 Element manager. Please enter valid Username and password.	1879
EMSG0039E Unable to Add or Modify DS8000 Element manager. The type of token submitted is not supported by the Storage Authentication Service.	1879
EMSG0040I The SMI-S provider Connection is about to be removed. Do you wish to continue?	1879
EMSG0041E The DS8000 Network server is unavailable.	1880
EMSG0042E You are not authorized to perform this action through the management console.	1880
EMSG0043E You cannot add an element manager for this DS8K as its software version is too low.	1880
EMSG0044E You cannot update the element manager for this DS8K as its software version is too low.	1880
GEN - General Spectrum Control messages	1880
GEN0001E command name(command arguments) FAILED.	1883
GEN0002E Unable to create temporary file file name.	1883
GEN0003E Unable to lock temporary file file name.	1883
GEN0004E accept() fails -- error message.	1884
GEN0005E Cannot get input stream from host <computer name>.	1884
GEN0006E Cannot get output stream to host <computer name>.	1884
GEN0007E Cannot deserialize from host <computer name>.	1884
GEN0008E Cannot read from host <computer name>.	1884
GEN0009E Cannot serialize to host <computer name>.	1885
GEN0010E Cannot write to host <computer name>.	1885
GEN0011E Cannot open object stream from host <computer name>.	1885
GEN0012E Socket to host <computer name> closed prematurely.	1885
GEN0013E Object read from host <computer name> is not Request -- it is.	1886
GEN0014E Object read from host <computer name> is not Response -- it is.	1886
GEN0015E Socket to host <computer name> -- class not found.	1886
GEN0016I Above error occurred sending Request (request type, request subtype).	1886
GEN0017I Above error occurred sending Request (request type, request subtype) (phase value).	1887
GEN0018I Response received from host <computer name>.	1887
GEN0019I Above error occurred responding to Request (request type, request subtype).	1887
GEN0020I Above error occurred reading data for Request (request type, request subtype).	1887
GEN0021E Error closing socket to host <computer name>.	1887
GEN0023E Cannot rename old file name to new file name.	1887
GEN0024I On socket to host <computer name>.	1888
GEN0025E Request(request type, request subtype) was incomplete, but response said complete.	1888
GEN0026E Request(request type, request subtype) was complete, but response said incomplete.	1888
GEN0027E Cannot send incomplete response to single-phase Request(request type, request subtype).	1888
GEN0028E Tried to read data of single-phase Request(request type, request subtype).	1888
GEN0029E Cannot open.	1889
GEN0030E Cannot create listener on port port number.	1889
GEN0031E Unroutable type-code.	1889
GEN0032E Unroutable sub-type.	1889
GEN0033W SoTimeout failed -- host <computer name>.	1890
GEN0035E Cannot seek to position on file name.	1890
GEN0036E Cannot create.	1890
GEN0037E Cannot create file name.	1890
GEN0038W Cannot delete file name.	1890
GEN0039E Unknown host <computer name>.	1891
GEN0040E Cannot connect to <computer name:port number>.	1891
GEN0041E System property <property name> is not defined.	1891
GEN0042E Error reading log-file file name.	1891

GEN0043E Cannot find SM/DMI header in low memory.	1891
GEN0044E Cannot seek to physical memory address address.	1892
GEN0045E No type-1 structure found.	1892
GEN0046E Obsolete type-1 structure (no UUID).	1892
GEN0047E Invalid or unset UUID.	1892
GEN0048E Cannot obtain system manufacturer.	1893
GEN0049E Cannot open log file file name.	1893
GEN0050E Unable to connect to database repository.	1893
GEN0052E Cannot open directory directory name.	1893
GEN0054E Error creating pipe.	1894
GEN0055E Error sussing pipe.	1894
GEN0056E Cannot perform operation <operation> on physical memory.	1894
GEN0057W Windows message message.	1894
GEN0058W Windows message message: description.	1894
GEN0067W Cannot write to pipe name.	1895
GEN0068E Cannot find binary module module name.	1895
GEN0069E Cannot find entry point entry point name.	1895
GEN0070E Cannot find privilege privilege name.	1895
GEN0071E Cannot open own token.	1895
GEN0072E Cannot assert privilege privilege name.	1896
GEN0073E Cannot look up network interfaces.	1896
GEN0074E No Ethernet cards found.	1896
GEN0075E GetTokenInformation() failed.	1896
GEN0076E Not super-user.	1897
GEN0077E SetHandleInformation() failed.	1897
GEN0078I Trying token-ring.	1897
GEN0079E socket() failed.	1897
GEN0080E Error looking up Ethernet card.	1897
GEN0081W No token-ring cards found.	1898
GEN0082W Error looking up token-ring card.	1898
GEN0083I Trying Ethernet.	1898
GEN0084I Hardware-ID obtained.	1898
GEN0096I PID = process identifier.	1898
GEN0097E Unable to retrieve hardware-ID.	1898
GEN0098E Error processing request from host <computer name>, user <user name>, for service <service name>, request(request type, request subtype).	1898
GEN0099W Warning processing request from host <computer name>, user <user name>, for service <service>, request(request type, request subtype).	1899
GEN0100E Error processing request from host <computer name>, for service <service>, request(request type, request subtype).	1899
GEN0101W Warning processing request from host <computer name>, for service <service>, request(request type, request subtype).	1899
GEN0102E Missing t_identifiers row. ID type id number.	1899
GEN0104E Not enough virtual memory.	1900
GEN0105E Unable to send internal job results to T-Time Schedule : scheduler, Job: job name Run Number: run number.	1900
GEN0106E Unable to obtain local hostname. hostname.	1900
GEN0107E Unable to obtain SNMP datagram socket.	1900
GEN0108E Unable to send SNMP trap datagram.	1900
GEN0109E Unable to deliver an e-mail to one or more of the following recipients: recipients	1901
GEN0110E Unable to read reply from SMTP server server name.	1901
GEN0111E Unable to connect to SMTP server server Authentication failed	1901
GEN0112E Unable to write to SMTP server server name.	1901
GEN0113E Unable to connect to SMTP server server name Unknown host.	1902
GEN0114E Unable to connect to SMTP server server name, port port number.	1902
GEN0115E SNMP server server name is unknown.	1902
GEN0125E Requested Report invalid: report name/report number.	1902
GEN0126W Unable to determine home directory.	1902
GEN0127W job creator.job name job messages will be logged to the Data Server log file.	1903
GEN0128E License key must be in format xxxxx-xxxxx-xxxxx-xxxxx.	1903
GEN0129E License key characters must be 0-9 or A-Z excluding E, I, O, U.	1903
GEN0130E License key checksum is incorrect - key is invalid.	1903
GEN0131E Unable to open file file name.	1903
GEN0132E Error parsing file name. Unrecognized section name: section name.	1904
GEN0133E Error reading file file name.	1904
GEN0134E Error parsing file file name Unrecognized token token name in section section name.	1904
GEN0135E Error - file file name appears to be truncated.	1904
GEN0136E Not enough agent licenses to run job.	1905
GEN0137E product name license has expired.	1905
GEN0139E product name is not installed on agent agent name.	1905
GEN0140E product name is not licensed for agent agent name.	1905
GEN0141E License key key is invalid.	1906
GEN0143E key license key is expired.	1906
GEN0148E Unknown product on computer name.	1906
GEN0149E product name is at release release level on agent name and at release release level on the server.	1906
GEN0151I product name vversion.release.modification.	1907
GEN0152E Error writing file name.	1907

GEN0153E Error serializing to file namefile name.	1907
GEN0154E Unable to retrieve cached report data Error reading file name.	1907
GEN0155E Unable to retrieve cached report data Error deserializing from file name.	1907
GEN0156E Unable to retrieve cached report data Premature end of file -- file name.	1908
GEN0157E Unable to retrieve cached report data Class class name not found restoring from file name.	1908
GEN0158E Unable to retrieve cached report data Object restored from class name is not file name; it's.	1908
GEN0159E Unable to retrieve cached report data Your request was interrupted. Processing terminated.	1908
GEN0160E Cached report data is no longer available on the server. Generate the report again. Cannot open.	1909
GEN0161E Error setting permissions on file name.	1909
GEN0162E Error retrieving FD flags.	1909
GEN0163E Error setting FD flags.	1909
GEN0164E Write timed out.	1910
GEN0165E poll() failed.	1910
GEN0166E write() failed.	1910
GEN0167E SO_SNDTIMEO failed.	1910
GEN0172E Error processing Request(request type, request subtype).	1910
GEN0173E Java Error in readObject(): error message.	1911
GEN0174E requestData is: class name.	1911
GEN0175E writeObject() failed writing: error message.	1911
GEN0176E responseData was: class name.	1911
GEN0177E Java Error in writeObject: error message.	1912
GEN0178E GuiReportReq report(report type, report subtype).	1912
GEN0179E GuiListReq listRequested(type).	1912
GEN0180W A license key exists for product name, but the software is not installed.	1912
GEN0181E Error - duplicate rows found for agent agent name.	1912
GEN0182I TsName: name, Manufacturer: manufacturer, HwID: hardware id.	1913
GEN0184E License key is invalid for this software release.	1913
GEN0197E Bad magic number.	1913
GEN0198I product component starting.	1913
GEN0199E The license key for this product has been deleted.	1913
GEN0200E Not enough licenses to license all requested agents.	1914
GEN0201W operating system : License not present or expired.	1914
GEN0222W Cannot find binary module module name.	1914
GEN0223W Cannot find entry point entry point name.	1914
GEN0304E Failed to read the stream header.	1914
GEN0305E Failed to decrypt input stream.	1915
GEN0306E Authentication failed from host:request type request subtype.	1915
GEN0307E Request not allowed from a cryptable stream request type,request subtype.	1915
GEN0308E Failed to write the stream header.	1915
GEN0309E Failed to re-authenticate.	1916
GEN0310E Failed to encrypt input stream.	1916
GEN0311E Authentication failed.	1916
GEN0329E Failed to authenticate the user <user ID>.	1916
GEN0332E Unable to connect to host hostname, port port.	1916
GEN0335E Unable to send a test email message to server name.	1917
GEN0400I Probe aborted.	1917
GEN0400E Probe completed with errors.	1917
GEN0401I Probe completed successfully.	1917
GEN0402W Probe completed with warnings.	1917
GEN0403E Unable to retrieve probe definition.	1917
GEN0404I Probe started.	1918
GEN0405E No computers or storage subsystems to probe.	1918
GEN0406E A probe of the computer or hypervisor computer name is already in progress.	1918
GEN0407E Probe of storage subsystem subsystem name (subsystem alias) already in progress.	1918
GEN0408W Probe of fabric fabric WWN already in progress.	1918
GEN0409E Batch Report job report name already in progress.	1919
GEN0410E Probe of switch switch WWN already in progress.	1919
GEN0557E An OS error occurred.	1919
GEN0558E An IO error occurred.	1919
GEN1030E product name is not licensed on computer computer name.	1919
GEN1034I GeneralException message text follows:error message.	1920
GEN1035W storage subsystem name is no longer monitored.	1920
GEN6013E OS Error errno: error description	1920
GEN6014W OS Error errno	1920
GEN6015E Unable to load JAVA.DLL	1920
GEN6016E Unable to find symbol module symbol in JAVA.DLL	1921
GEN6017E Executable file is null	1921
GEN6018E stat(file path) failed.	1921
GEN6019E chmod(file path) failed.	1921
GEN5001W Not all statistics could be saved due to unlicensed computer(s). See TPCD log for details.	1921
GEN7111E Unable to retrieve data from the repository database.	1922
GEN7112E Failed to retrieve e-mail report from (report URL).	1922

GEN7113E Failed to retrieve e-mail report from the web server.	1922
GEN7114E The user and password combination for the outgoing e-mail server is invalid.	1922
GEN7115E Can't send the attached "reportName" report with the ID reportId. The size of the attachment (totalSize KB) exceeds the maximum size for attachments that was set on your email server. Learn More link	1923
GEN7116E The email address that was configured for replies to alert notifications and reports, reply_to_address_value, was not accepted by the email server, server.	1923
GPC - Performance User Interface messages	1923
GPC000001E One or more of the selected profiles is a system-defined profile and cannot be deleted. Deselect profile and try again.	1923
GPC000002E The selected profile profile is a system-defined profile and cannot be updated.	1924
GPC000003E There is no performance data collected for any of the storage subsystems, create workload profile Wizard can not continue.	1924
GPC000150W Analyzing all the volume performance data may take a long time to complete.	1924
GPC000200E Based on the current choices, no devices can be included in the analysis.	1924
GPC000201E There is no ESS performance data collected, Volume Performance Advisor Wizard can not continue.	1924
GPC000350I Volume size may round up to multiple of 100 MB.	1924
GPC00400E Subsystem type for type identifier does not match the subsystem type in selected list.	1925
GPC00401E Duration value is not specified.	1925
GPC00800E The value entered for the number of rows is incorrect.	1925
GPC00801E The selected metrics do not have same unit type. Charts only display metrics with the same unit type.	1925
GPC00802E The selected charting metrics do not have same chart type. Chart and History Chart options cannot be mixed.	1925
GPC00803E Performance reports can only be saved with 300 or less explicitly specified components. To save the report either select all components or select no more than 300.	1925
GPC00804E Please enter a relative time higher than zero.	1926
GPC00950I There is no data to be charted.	1926
GPC00951I No metrics were selected.	1926
GPC00952I No drill-up available for this component.	1926
GPC00953I None of the selected metrics apply for all the report components.	1926
GPC00954I None of the selected metrics apply for all selected components.	1926
GPC00955I None of the selected metrics apply for all the report components or for all selected components.	1927
GPC00956I Subsystem port constraint violations do not support the affected volumes report.	1927
GPC00957I The maximum limit of max no of rows displayable rows for a performance report has been reached. If a larger report is required rerun the report as a batch report, alternatively redefine the report to return less data. The amount of report data can be reduced by using a smaller time window, working with less components, applying filters when defining the report or using the aggregated hourly or daily report data rather than the sample data.	1927
GPC00958I No drill-down available for this component.	1927
GPC50000E Fabric type for fabric type identifier does not match the fabric type in selected list.	1927
HWNAS - Agentless Server messages	1928
HWNAS0001I Successfully created server server.	1928
HWNAS0002I Successfully deleted server server.	1928
HWNAS0003E The host name or IP address {0} is not valid.	1928
HWNAS0004E Cannot add port portWWPN because it belongs to server serverName.	1928
HWNAS0005E Cannot add port portWWPN because it belongs to switch switchName.	1929
HWNAS0006E Cannot add port portWWPN because it belongs to storage system storageSystemName.	1929
HWNAS0007W Server serverName was not created because it exists already.	1929
HWNAS0008I Successfully created mergeServerName server by merging numberOfServers servers.	1929
HWNAS0009I Successfully separated serverName server into numberOfServers individual servers.	1929
HWNAS0010E The serverId agentless server that you selected does not exist.	1930
HWNAS0011I You cannot separate the serverName agentless server because it is not based on storage system host connections.	1930
HWNAS0012I You cannot separate the serverName agentless server because it is already defined on the smallest possible separation boundary.	1930
HWNAS0013I You cannot merge the selected agentless servers into the serverName agentless server because they are not all based on storage system host connections.	1930
HWNAU - Single sign-on service messages	1930
HWNAU0001E A connection with the IBM Spectrum Control Device Server, (Device Server IP, could not be established for authentication.	1931
HWNAU0002E A connection with the LDAP or Active Directory server, (LDAP or Active Directory Server IP) , could not be made for authentication.	1931
HWNAU0003E Authentication of the Single Sign-On token failed. Provide your username and password to attempt a re-authentication.	1931
HWNAU0004E The Single Sign-On token has expired. To re-authenticate the token, please enter your user name / password.	1931
HWNAU0005E Creation of the Single Sign-On token failed due to an username that is not valid. Enter your username and password and try again.	1932
HWNAU0006E Creation of the Single Sign-On token failed due to a password that is not valid. Enter your username and password and try again.	1932
HWNAU0007E Authentication failed due to an username or password that is not valid. Enter your username and password and try again.	1932
HWNAU0008I Single Sign On Service started successfully.	1932
HWNAU0009I The Single Sign On Service has shutdown.	1932
HWNAU0010E An error occurred retrieving the Single Sign-On token from the private credentials.	1933
HWNAU0011E An error occurred retrieving the Single Sign-On token from the public credentials.	1933
HWNAU0012E An error occurred when attempting to decode the authentication token.	1933
HWNAU0013E An error occurred when attempting to encode the authentication token.	1933
HWNAU0014E An error occurred while translating the user's credentials into a Single Sign-On token.	1933
HWNAU0015E An unknown error occurred while authenticating to the WebSphere login module.	1933
HWNAU0016E An error occurred while registering SsoConfigChangeListener with TIP.	1933
HWNAU0017E An error occurred while unregistering SsoConfigChangeListener from TIP.	1934
HWNAU0018E The Web server appears to be down and cannot be used for authentication. It is still possible to perform OS user authentication against the device server, however since the Web server is down the IBM Spectrum Control functionality will be limited. Among the limitations is the inability to perform SSO to other applications that rely on the presence of a lightweight third party authentication token. To proceed enter a local OS username with administrative privileges and password.	1934
HWNAU0019E An unknown error occurred while authenticating with Web server.	1934
HWNCA - Storage multiple access control messages	1934

HWNCA0001E The required parameter_name parameter is missing.	1935
HWNCA0002E The delivery unit type delivery_unit_type is not supported.	1935
HWNCA0003E The delivery unit type, delivery_unit_type, does not match the capacity pool type, capacity_pool_type.	1936
HWNCA0004E The service class service_class_name is not supported by the capacity pool.	1936
HWNCA0005E IBM Spectrum Control cannot provide the requested delivery_unit_size GiB of capacity that also satisfies the requirements of the service class.	1936
HWNCA0006E A delivery unit delivery_unit_name and service instance ID service_instance_ID already exists.	1936
HWNCA0007I Normal isolated file storage.	1936
HWNCA0008I Enhanced isolated file storage.	1937
HWNCA0009E The capacity pool with ID capacity_pool_id is still in use and cannot be deleted.	1937
HWNCA0010E The service class service_class_name has no recipe for deleting delivery unit delivery_unit_name.	1937
HWNCA0011E The service class service_class_name cannot modify the delivery unit delivery_unit_name.	1937
HWNCA0012E The capacity of the delivery unit with ID delivery_unit_id cannot be modified after delivery unit is created.	1937
HWNCA0013E Delivery unit with ID delivery_unit_id cannot be modified or deleted because it is in the processing state.	1938
HWNCA0014E Delivery unit with ID delivery_unit_id cannot be modified because it has a completion state error.	1938
HWNCA0015E The principal principal_name was used in multiple CIFS access control list (ACL) entries.	1938
HWNCA0016I A standard object can handle object-type delivery units, such as CDMI containers.	1938
HWNCA0017I You can use the StandardBlock service class to create, modify, and delete delivery units that are block based, such as volume containers.	1938
HWNCA0018I High Availability, Sequential.	1939
HWNCA0019I High Availability, Short response time, Transactional.	1939
HWNCA0020E No provisioning profile was found for recipe name.	1939
HWNCA0021E The WWPN WWPN does not belong to a server or hypervisor managed by IBM Spectrum Control.	1939
HWNCA0022E Device Selection failed with the following error: Planner Message	1939
HWNCA0023E Creation of DeliveryUnit failed with the following error: Planner Message	1939
HWNCA0025I Remote Caching allows the data written to one delivery unit to be pushed to another remote located delivery unit.	1940
HWNCA0026E Deletion of DeliveryUnit failed with the following error: Planner Message	1940
HWNCA0027I Highest-performing storage for mission-critical applications.	1940
HWNCA0028I High-performing storage for applications in production.	1940
HWNCA0029I Standard storage for non-mission-critical applications.	1940
HWNCA0030E No host port WWPNs were found for server host name.	1940
HWNCA0031I The standard root object class enforces the use of dedicated filesets for associated object-type delivery units, such as CDMI containers.	1941
HWNCA0032E A Storage Resource agent is not up and running for the host name server.	1941
HWNCA0033E A service class with name service_class_name does not exist for the specified delivery unit type.	1941
HWNCA0034E IBM Spectrum Control is unable to create or modify the service class because one or more of the capacity pools to which it was restricted were removed.	1941
HWNCA0035E Delivery unit with ID delivery_unit_id cannot be modified or deleted because it has a completion status of planned or change_planned.	1942
HWNCA0036E More than numResults results found. Modify your search pattern.	1942
HWNCA0037I Start the creation of an object container with name containerName.	1942
HWNCA0038I The creation of the object container with name containerName is finished.	1942
HWNCA0039I Start the deletion of an object container with name containerName.	1942
HWNCA0040I The deletion of the object container with name containerName is finished.	1943
HWNCA0041I Start the modification of an object container with name containerName.	1943
HWNCA0042I The modification of the object container with name containerName is finished.	1943
HWNCA0043I This task was already deleted.	1943
HWNCA0044E The requested operation requires SMAC API client version version or later.	1943
HWNCA0045E The plan with scheduleID scheduleID can not be deleted because the delivery unit creation is already processing state.	1944
HWNCA0046E The specified value deletionType is not a valid plan deletion type.	1944
HWNCA0047E Host with host name host is not known by IBM Spectrum Control.	1944
HWNCA0048E Filer Storage ID filerId for given fileSystem is not found.	1944
HWNCA0049W IBM Spectrum Control did not change any Fibre Channel zones. After the provisioning operation completes, verify the Fibre Channel connectivity between any involved host and the storage subsystem that contains the storage volume. If necessary, change the zoning configuration.	1944
HWNCA0050W IBM Spectrum Control did not enable the multipath policy for the host or hosts to which the volume was reassigned. You must enable the multipath policy after the provisioning operation completes.	1945
HWNCA0051E Incorrect volume name volumeName. The name can contain letters, numbers, dashes (-) and underscores (_). However, the name cannot start with a number, dash, or the reserved words vdisk or volume.	1945
HWNCA0052E One or more of the requested volume names are too long for the following storage systems: StorageSubsystemsList. The maximum volume name length that is allowed is MaxAllowedNameLength characters, and a volume name of LargestNameLength characters was requested.	1945
HWNCA0053E The service class scName does not support provisioning from any available storage.	1945
HWNCA0054E The volume volumeName already exists on the storage system storageSystemName. Request a different volume name.	1946
HWNCA0055E The share shareName already exists on the storage system storageSystemName. Request a different share name.	1946
HWNCA0056W Zoning cannot be configured because the connected fabrics are not monitored.	1946
HWNCA0057W The multipath policy cannot be set on the following hypervisors: HostList.	1946
HWNCA0058W The multipath policy cannot be set on the following agentless servers: HostList.	1946
HWNCA0059W The multipath policy cannot be set on the following hosts: HostList.	1947
HWNCA0060E Storage volume can not be assigned to servers that are running different operating systems.	1947
HWNCA0061E The required resources cannot be reserved because too many provisioning tasks are currently being created. Try again in a few minutes.	1947
HWNCA0062E The operation did not complete because the specified service class or capacity pool no longer exists.	1947
HWNCA0063E The service class ScName that is specified by this provisioning task does not exist.	1947
HWNCA0064E The current SMAC API client is version current_version. The requested operation requires SMAC API client version required_version or later.	1948
HWNCA0065E The delivery unit specified has snapshots and could not be deleted.	1948
HWNCA0066E The WWPN WWPN already belongs to a server or hypervisor managed by IBM Spectrum Control. Was expected an unknown WWPN based on previous Hosts or WWPNs passed in.	1948

HWNCA0067E The host host already belongs to a server or hypervisor managed by IBM Spectrum Control. Was expected an unknown Host based on previous Hosts or WWPNS passed in.	1948
HWNCA0068E The list of hosts was expected to contain just one or no element because the WWPNS passed in were all unknown. The size of the host list is size	1949
HWNCA0069E No unknown WWPNS were specified for this provisioning operation but an unknown Host was specified.	1949
HWNCA0070E The expansion of the capacity is not supported for the delivery unit with ID delivery_unit_id.	1949
HWNCA0071E The reduction of the capacity is not supported for the delivery unit with ID delivery_unit_id.	1949
HWNDA - Data Manager API messages	1949
HWNDA0001I Operation Name of the operation processed successfully.	1951
HWNDA0002E Mandatory parameter Name of the mandatory parameter which is missing missing	1951
HWNDA0003E Invalid parameter Name of the parameter which was invalid	1951
HWNDA0004E An internal error occurred.	1951
HWNDA0005E The server encountered an error when it was accessing the database.	1951
HWNDA0006E The name provided while creating a new group is already in use.	1952
HWNDA0007E An external key could not be identified for the provided type The constant integer type of the Group element and id The unique integer database ID of the Group element.	1952
HWNDA0008E The specified attribute invalid attribute name is not a valid attribute.	1952
HWNDA0009E An internal ID could not be identified for the provided type The constant integer type of the Group element and external key The unique external key of the Group element.	1952
HWNDA0010I The following elements are already members of the group The group: The keys of the elements.	1953
HWNDA0011I The following elements are not members of the The group Group and cannot be removed: The element key.	1953
HWNDA0012E Adding a Group with the name Name of the proposed new member to the Name of the parent group Group would create a circular relationship that is not allowed.	1953
HWNDA0013E The input parameter value input parameter value for input input parameter name exceeds the maximum allowable length of number of allowable characters characters.	1953
HWNDA0014E The provided Group attribute value Group attribute value for the Group attribute name Group attribute contains invalid characters. The following characters are not allowed, \\\\:.*?>< ."	1953
HWNDA0015E You are not the original creator of the provided Group name Group name.	1954
HWNDA0016E The provided Tiering Policy name Tiering Policy name is not unique.	1954
HWNDA0017E The provided Group Group name or ID does not exist.	1954
HWNDA0018E The provided Tiering Policy name Tiering Policy name does not exist.	1954
HWNDA0019E The provided candidate and destination Group names, Group name, cannot be the same.	1954
HWNDA0020E The provided condition condition type is not valid.	1955
HWNDA0021E The provided operand operand type is not valid.	1955
HWNDA0022E The provided condition condition type is either already applied to this tiering policy or conflicts with an existing condition, existing condition type	1955
HWNDA0023E The requested priority value priority value is invalid.	1955
HWNDA0024E The specified Group name Group name is not unique.	1955
HWNDA0025E Cannot add the specified resource because the resource type, element type, is not supported as a child of the group.	1956
HWNDA0026E Cannot add the specified group, Group name, because the group type, type, is not supported as a child of the application.	1956
HWNDA0027E The first option specified in the file must be -appgroupname.	1956
HWNDA0028E The argument of the option option is missing at or before line Line Number : Line	1956
HWNDA0029E Both option1 and option2 were specified at or before line Line Number : Line	1956
HWNDA0030E The option option is missing at or before line Line Number : Line	1957
HWNDA0031E Neither option1 nor option2 was specified at or before line Line Number : Line	1957
HWNDA0032E Invalid number of parameters for option option at or before line Line Number : Line	1957
HWNDA0033E Incomplete options sequence before end of file.	1957
HWNDA0034E Invalid option option at line Line Number : Line	1957
HWNDA0035E Invalid resource type type at line Line Number : Line	1958
HWNDA0036E Invalid sequence of options at or before line Line Number : Line	1958
HWNDA0037E Syntax error, quote sequence not properly closed at line Line Number : Line	1958
HWNDA0038E Option option is not allowed for resource type type at or before line Line Number : Line	1958
HWNDA0039E The input data for modifying the application groups is missing.	1958
HWNDA0040E An invalid element was encountered in the input data.	1958
HWNDA0041E The application group name is missing from the input data.	1959
HWNDA0042E The operation is missing from the input data.	1959
HWNDA0043E The resource type is missing from the input data.	1959
HWNDA0044E The server name is missing from the input data.	1959
HWNDA0045E The device name is missing from the input data.	1959
HWNDA0046E Invalid values were specified for the server, device or cluster names in the input data.	1960
HWNDA0047E Member names and the tags were specified for the same operation.	1960
HWNDA0048W The following entities were not found: Entities.	1960
HWNDA0049W No entities were found for the tags: Tags	1960
HWNDA0050E The member names or tags were not specified for the operation.	1960
HWNDA0051W The group Name of the proposed new member cannot be added to itself.	1960
HWNDA0052W The group Name of the proposed new member cannot be added to the Name of the parent group group because it creates a circular relationship that is not allowed.	1961
HWNDA0053W The group Name of the group contains child groups and cannot be deleted. The child groups must be deleted before the parent group can be deleted.	1961
HWNDA0054E The filter mask that was used to create or edit a group filter is currently being used.	1961
HWNDA0055E The argument argument for the parameter parameter on line Line Number : Line is invalid.	1961
HWNDA0056E The first option specified in the file must be -id.	1961
HWNDA0057E The specified Group Group name is not of type type.	1962
HWNDA0058W These groups have same names as existing members of the group The group: The keys of the elements. They were not added to the group.	1962
HWNDA0059E The specified tag key Tag key is not valid.	1962

HWNDA0060W These group members are also members of another group: The groups. They are not deleted.	1962
HWNDA0061W These group members cannot be moved up one level in hierarchy due to name conflicts: The groups. The group is not deleted.	1962
HWNDA0062E The specified Group Group name is not a department group.	1963
HWNDA0063W These group members cannot be moved as top level groups in hierarchy due to name conflicts: The groups. No group members were removed from the group.	1963
HWNDA0064E An application with the same name already exists.	1963
HWNDA0065E A department with the same name already exists.	1963
HWNDA0066E Invalid values were specified for the device, cluster or file system names in the input data.	1963
HWNDA0067E Resources of type Type of resource cannot be added to an application or removed from an application using tags.	1964
HWNDA0068E The application cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.	1964
HWNDA0069E The department cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.	1964
HWNDA0070W The File Systems: {0} were not added to the application {1} because they are NAS file systems.	1964
HWNDA0071E Member IDs and the tags were specified for the same operation.	1964
HWNDA0072E Member IDs should be specified in the input data for resources of type appgroup.	1965
HWNDA0073E A general group with the same name already exists.	1965
HWNDA0074E The general group cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.	1965
HWNDA0075E A dashboard with the same name already exists.	1965
HWNDA0076E The dashboard group cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.	1965
HWNDA0077E A policy group with the same name already exists.	1966
HWNDA0078I The Name of the policy group policy group was removed.	1966
HWNEM - Element manager messages	1966
HWNEM0001E The element manager management service failed to obtain a database connection.	1967
HWNEM0002E An error occurred while attempting to add element manager information to the database.	1967
HWNEM0003E An error occurred while attempting to update element manager information in the database.	1968
HWNEM0004E An error occurred while attempting to remove element manager information from the database.	1968
HWNEM0005E An error occurred while attempting to obtain element manager information from the database.	1968
HWNEM0006E Failed to transmit request to Data Server to initiate SMI-S provider discovery job.	1968
HWNEM0007E An error occurred while attempting to obtain Data Server information from the database.	1969
HWNEM0008E Data Server information was not found in the database. SMI-S provider discovery could not be scheduled.	1969
HWNEM0009E An error occurred while attempting to obtain the SMI-S provider URL associated with an element manager from the database.	1969
HWNEM0010E An error occurred while attempting to obtain element manager credentials from the database.	1970
HWNEM0011E An error occurred while attempting to remove element manager credentials from the database.	1970
HWNEM0012E An error occurred while attempting to store element manager credentials in the database.	1970
HWNEM0013E Encountered element manager with malformed URL (URL).	1970
HWNEM0014E Encountered element manager with URL containing hostname that could not be resolved by DNS (URL).	1971
HWNEM0015E Failed to clone element manager.	1971
HWNEM0016E Failed to encrypt element manager password.	1971
HWNEM0017E Failed to decrypt element manager password.	1971
HWNEM0018E Failed to send request to Data Server (dataServerHost:dataServerPort).	1972
HWNEM0019I Attempting to schedule discovery on Data Server for number SMI-S provider(s) (SMI-S provider URLs)...	1972
HWNEM0020I Successfully scheduled discovery on Data Server for number SMI-S provider(s).	1972
HWNEM0021I SMI-S provider discovery was not scheduled on Data Server. No SMI-S providers are associated with the specified set of element managers.	1972
HWNEM0022E Failed to authenticate with ESSNI server associated with element manager at URL using ESSNI user ID ESSNI user ID.	1973
HWNEM0023E The ESSNI server associated with the element manager at URL is not available.	1973
HWNEM0026E An error occurred while attempting to retrieve version information from the ESSNI server associated with the element manager at URL.	1973
HWNEM0029E Failed to set element manager credentials. Specified Element manager does not exist.	1973
HWNEM0100E The element manager's URL is not in the correct format.	1974
HWNEM0101E Change the default administrator password now to avoid security conflicts.	1974
HWNEM0102E The element manager already exists.	1974
HWNEM0103E A problem occurred adding the element manager to IBM Spectrum Control.	1974
HWNEM0104E Login to the element manager failed. Check the user credentials.	1975
HWNEM0105E The request contained data in an unexpected format	1975
HWNEM0106E The request did not contain the expected information.	1975
HWNEM0107E The request contained an unsupported action.	1975
HWNEM0108E Connection test to SMI-S provider URL FAILED due to status code.	1975
HWNEM0109E Connection test to Element Manager failed.	1976
HWNEM0111E A problem occurred locating the element manager in IBM Spectrum Control.	1976
HWNEM0112E A problem occurred updating the element manager to IBM Spectrum Control.	1976
HWNEM0113E A problem occurred removing the SMI-S provider from the element manager.	1976
HWNEM0114E A problem occurred locating the SMI-S provider associated with the element manager.	1977
HWNEM0115E Unable to establish an https connection to the element manager.	1977
HWNEM0116E A problem occurred removing the element manager from IBM Spectrum Control.	1977
HWNEM0117E A problem occurred testing the connection to the element manager.	1977
HWNEM0118I Connection test to the element manager element manager passed.	1978
HWNEM0119E Connection test to the element manager element manager failed.	1978
HWNEM0120E IBM Spectrum Control is unable to communicate with the element manager, the problem could be the element manager is not running, a network communication error or user credentials stored in IBM Spectrum Control are incorrect.	1978
HWNEM0121I Connection test to the SMI-S provider SMI-S provider passed.	1978
HWNEM0122I Connection test to the SMI-S provider SMI-S provider failed.	1979
HWNEM0123E Before executing the action the user must select an element manager from the table.	1979
HWNEM0124E There is no SMI-S provider associated with the selected element manager.	1979

HWNEM0125I The element manager is about to be removed. Once removed the element manager will not be accessible from IBM Spectrum Control. Do you wish to continue.	1979
HWNEM0126I The element manager is about to be removed. Once removed the element manager will not be accessible from IBM Spectrum Control. Do you wish to continue.	1979
HWNEM0127E An internal processing error occurred while servicing the last request.	1980
HWNEM0128E The element manager is not available. Ensure that the element manager's credentials are defined and up to date.	1980
HWNEM0129E An unexpected error occurred changing the element manager's default password.	1980
HWNEM0130E SMI-S provider connection was added to the element manager, however the SMI-S provider discovery job failed to launch.	1980
HWNEM0131I The DS8000 SMI-S provider has been added successfully. IBM Spectrum Control has started the discovery job for the DS8000 storage subsystem managed by this SMI-S provider. To check the status of the jobs, go to the IBM Spectrum Control perspective and check the following navigation tree nodes: Administrative Services -> Discovery -> CIMOM	1981
HWNEM0132I SMI-S provider connection was removed from element manager.	1981
HWNEM0133E Internal error occurred, the element manager info for launching the element manager could not be located in the IBM Spectrum Control DB.	1981
HWNEM0134E The user session data is no longer available. To continue restart the IBM Spectrum Control GUI.	1981
HWNEM0135E Unable to contact the device server. It appears to be down.	1982
HWNEM0136I The DS8000 Element Manager at IP Address has been added successfully.	1982
HWNEM0137I service.method IBM Spectrum Control User IBM Spectrum Control user launched DS8000 Element Manager IP Address under the alias of DS8000 user name.	1982
HWNEM0138I The DS8000 SMI-S provider has been modified successfully. IBM Spectrum Control has started the discovery job for the DS8000 storage subsystem managed by this SMI-S provider. To check the status of the jobs, go to the IBM Spectrum Control perspective and check the following navigation tree nodes: Administrative Services -> Data Sources -> Discovery -> CIMOM	1982
HWNEM0139I The DS8000 SMI-S provider has been modified successfully.	1982
HWNEM0140W The element manager GUI cannot be accessed because no username and password have been specified by the currently logged-in IBM Spectrum Control user. Element manager credentials are managed on a per-IBM Spectrum Control user basis. The element manager GUI will appear after you provide the correct username and password information.	1983
HWNEM0141E The length of the SMI-S provider's text description is too long. The description should be 255 characters or less.	1983
HWNEM0142E An element manager named manager.name already exists.	1983
HWNEM0143E The DS8000 Element Manager Console is only accessible from IBM Spectrum Control. To access it, open IBM Spectrum Control and switch to the DS8000 Element Manager perspective.	1983
HWNFNS - File system monitor messages	1984
HWNFNS0001E The file system monitor tool does not recognize the command: command.	1984
HWNFNS0002E The file system monitor tool is unable to read its property file.	1984
HWNFNS0003W The FileSystems property value property_value has an invalid format.	1984
HWNFNS0004E The file system monitor tool was unable to initialize access to the IBM Spectrum Control database repository.	1984
HWNFNS0005E The probe of device_name ended with the error_code error code.	1985
HWNFNS0006E The FileSystems property is missing in the fsmon.properties file.	1985
HWNFNS0007E Data cannot be retrieved from the database repository.	1985
HWNFNS0008E Data about IBM SONAS devices or file systems cannot be retrieved from the database repository.	1985
HWNFNS0009W Capacity data cannot be collected for the file_system file system.	1985
HWNFNS0010W No matching file system found in the database repository for file_system.	1986
HWNLML - Planner manager messages	1986
HWNLML0700E Storage pool recommendations cannot be given because a storage subsystem is not selected as input.	1990
HWNLML0701E The selected planner cannot occur because a storage subsystem is not selected as input.	1990
HWNLML0702E The selected planner cannot occur because a computer is not selected as input.	1990
HWNLML0703E The selected planner(s) cannot occur because a computer and/or volume is not selected as input.	1991
HWNLML0704E The selected planner cannot occur because a volume is not selected as input.	1991
HWNLML0705E Storage pool recommendation can not be done because there is already a storage volume selected as input.	1991
HWNLML0706E The selected planner cannot occur because there is already a volume selected as input.	1991
HWNLML0707E The selected planner(s) cannot occur because a subsystem or computer is not selected as input.	1991
HWNLML0708E The total required space (total space MB) exceeds the total available space (total space MB) considering all planner inputs. This could be due to insufficient space in the selected storage subsystem's controller(s) or pool(s), or the selected RAID level does not have enough free space or in case of SVC there is insufficient IO group memory configuration.	1992
HWNLML0709E The storage cannot be assigned to the specified hosts due to insufficient LUN addresses.	1992
HWNLML0710E The specified total size to be allocated (total space MB) is invalid. It must be positive and an integral multiple of 100.	1992
HWNLML0711E The specified minimum volume size (minimum size MB) is invalid. It must be positive and an integral multiple of 100.	1992
HWNLML0712E The specified maximum volume size (maximum size MB) is invalid. It must be positive and an integral multiple of 100.	1992
HWNLML0713E The total capacity requested cannot be obtained. The allowed minimum size is {0} GB and the allowed maximum size is {1} GB for the given subsystem(s). The total size is {2} GB.	1993
HWNLML0714E Host Host does not have multipath support because it only has one Fibre Channel port.	1993
HWNLML0715E There are no common fabrics that have the minimum number of required paths between the selected servers and the managed storage systems.	1993
HWNLML0716E There needs to be at least two common fabrics between host Host and storage subsystem Subsystem in order to use the redundant fabric option.	1993
HWNLML0718E The number of paths specified was Paths, but the redundant fabric option requires the number of paths to be an even number of paths.	1994
HWNLML0719E Host host does not have a Host Bus Adapter (HBA) installed. Please select a host with a HBA installed and try again.	1994
HWNLML0720E The Plan failed to generate due to storage subsystem Subsystem not having any volumes identified for Planning use. Please select another storage subsystem and try again.	1994
HWNLML0721E A supported multipath driver is not installed on host Host.	1994
HWNLML0722E There is an insufficient number of Fibre Channel paths between host Host and storage subsystem Subsystem. Requested Paths paths were requested but there are only Possible Paths paths available.	1994
HWNLML0723E The number of paths specified was Paths, but the redundant fabric option requires the minimum number of paths for virtual disks to be four paths or more.	1995
HWNLML0724E The number of zones in fabric fabric will be number of zones. This is larger than the max zones maximum number of zones specified.	1995
HWNLML0725E Storage pool data does not exist for storage subsystem subsystem. This is either due to not running a storage subsystem probe or not having any fixed block formatted storage pools on the subsystem.	1995

HWNLM0726E Performance data does not exist for input storage subsystem(s) for the given date range. Either run an IBM Spectrum Control Performance Monitor against the given storage subsystem(s) or under the Capacity Planner, select the 'Space Only' Workload Profile option.	1995
HWNLM0728E An unexpected internal error occurred. Please contact IBM customer technical support.	1996
HWNLM0729E The resulting SAN Planner actions include creating a zoned set on McData fabric fabric WWN. Since there is already an active zone set, please select the 'Use active zone set' option under the Zone Planner and try again.	1996
HWNLM0730E SDD version 1.6.2.3 installed on HP host host is not supported by IBM Spectrum Control.	1996
HWNLM0731E The computer probe for computer host was incomplete. Please attempt another computer probe and try again.	1996
HWNLM0732E DM-Multipath installed on Linux host host does not support the selected Multipath mode.	1997
HWNLM0734E The total required IO Group mirroring memory (total memory KB) for vdisk mirroring creation exceeds total available mirroring memory (total memory KB) available.	1997
HWNLM0735E Could not find any subsystem for planning. It is possible that the given subsystems are not detectable.	1997
HWNLM0736E The primary, secondary or tertiary storage subsystem subsystem that you specified is not registered with the IBM Spectrum Control-R server.	1997
HWNLM0737E Storage Subsystem subsystem is not a supported type for replication operations.	1997
HWNLM0738E Replication session with a combination of virtualized and non-virtualized storage volumes in not supported.	1998
HWNLM0739E Replication session type session is not allowed for the supplied subsystem subsystem of typesubsystemtype.	1998
HWNLM0740E Limit reached : Number of copy pairs for Replication session type session for the supplied subsystem subsystem.	1998
HWNLM0741E Replication session type session is not supported.	1998
HWNLM0742E Format of LSS Property File filepath is invalid.	1998
HWNLM0743E Namespace in specified LSS is not available.	1999
HWNLM0744E Unable to find suitable placement for replication storage volumes. plannermsg	1999
HWNLM0745E Secondary SRG secsrc is either empty or does not contain any valid elements for replication related resource provisioning.	1999
HWNLM0746E Tertiary SRG tersrc is either empty or does not contain any valid elements for replication related resource provisioning.	1999
HWNLM0747I Please ensure that proper Replication license and device feature codes are enabled, otherwise the plan execution will fail.	1999
HWNLM0748E Replication Manager is not installed.	2000
HWNLM0749E Replication Planner Internal Error.	2000
HWNLM0750E Replication session with the specified name (sesname) already exists. Please use a different name.	2000
HWNLM0751E No storage system resource can satisfy the provisioning requirements.	2000
HWNLM0752E The subsystem(s) provided in the input do not satisfy the given Thin Provisioning criteria.	2000
HWNLM0753E The subsystem(s) provided in the input do not allow provisioning on volumes or no subsystems were selected.	2001
HWNLM0754E The subsystem(s) provided in the input do not allow provisioning on virtual disks or no subsystems were selected.	2001
HWNLM0755E The selected virtual disk(s) cannot be added in a plan for volumes.	2001
HWNLM0756E The selected volume(s) cannot be added in a plan for virtual disks.	2001
HWNLM0757I The volume size has been slightly adjusted to meet subsystem requirement.	2001
HWNLM0758I If a FlashCopy source has multiple targets, an IncrementalFlashCopy relationship can be established with one and only one target.	2002
HWNLM0759I IncrementalFlashCopy is not available with FlashCopy SE.	2002
HWNLM0760W Input Volume is already in source role for 12 Flash Copy Sessions, ID = volumeID, Volume Name = volumeName.	2002
HWNLM0761W Input Volume is already in a target role of Flash Copy Session(s), ID = volumeID, Volume Name = volumeName.	2002
HWNLM0762W Input Volume is already in a target role of Continuous Copy Session(s), ID = volumeID, Volume Name = volumeName.	2002
HWNLM0763W Input Volume is already in a target role of Flash Copy Session(s).	2002
HWNLM0764I Ensure connectivity between source and target (direct or through fabric).	2003
HWNLM0765W For TSE volumes, please ensure that the Repository Capacity is configured and available on the pool poolName.	2003
HWNLM0766W The storage subsystem ssName is unacceptable due to -- reason.	2003
HWNLM0767I The storage subsystem ssName is a valid candidate subsystem, thus it is considered during planning.	2003
HWNLM0768W The storage pool poolName is unacceptable due to -- reason.	2003
HWNLM0769I The storage pool poolName is a valid candidate storage pool, thus it is considered during planning.	2003
HWNLM0770W The SVC ssName is unacceptable due to -- reason.	2003
HWNLM0771I The SVC ssName is a valid candidate subsystem, thus it is considered during planning.	2004
HWNLM0772W The mdiskgroup mdiskGroupName is unacceptable due to -- reason.	2004
HWNLM0773I The mdiskgroup mdiskGroupName is a valid candidate mdiskgroup, thus it is considered during planning.	2004
HWNLM0774W The iogroup ioGroupName is unacceptable due to -- reason.	2004
HWNLM0775I The iogroup ioGroupName is a valid candidate iogroup, thus it is considered during planning.	2004
HWNLM0776E Unable to plan for volumeName due to reaching the maximum limit of volumes on the LSS(es) within storage pool storagePool on storage subsystem subsystem.	2004
HWNLM0777E Replication session with the specified name (sesname) associated with SRG (srgrname) has required replication and does not require extension.	2005
HWNLM0778I The volume/vdisk name may be different during plan execution based on the name availability and/or subsystem limitation.	2005
HWNLM0779E Unable to recommend plan due to reaching the maximum limit of volumes on the LSS(es) within storage pool(s) storagePool on storage subsystem(s) subsystem.	2005
HWNLM0780E Unable to recommend plan since the LSS range specified in the LSSRange.properties file is not valid for volumeName on storage subsystem subsystem.	2005
HWNLM0781E No cluster partnership exists between source subsystem srcSS and target subsystem targetSS.	2006
HWNLM0782E No connectivity path exists between source subsystem srcSS and target subsystem targetSS.	2006
HWNLM0783E The selected input volume volumeName is missing from the storage subsystem.	2006
HWNLM0784E The selected input virtual disk virtualDiskName is missing from the SVC.	2006
HWNLM0785E Provisioning with replication can not be achieved with Extent Space Efficient(ESE) or Track Space Efficient(TSE) volumes. Extent Space Efficient volumes are not allowed in copy sets and Track Space Efficient volumes are valid only in the target role of a FlashCopy session or the Journal role of a Global Mirror or Metro Global Mirror session.	2006
HWNLM0786E Replication can not be extended to Extent Space Efficient(ESE) or Track Space Efficient(TSE) source volumes. Extent Space Efficient volumes are not allowed in copy sets and Track Space Efficient volumes are valid only in the target role of a FlashCopy session or the Journal role of a Global Mirror or Metro Global Mirror session. One or more selected input volumes are either ESE or TSE : volumes.	2007
HWNLM0787E vDisk Mirroring is already enabled for the input vdisk volumes.	2007
HWNLM0788I Fabric selection is based on planner selection of the host port(s) and the subsystem port(s) and not by the fabric(s) selected by the user.	2007
HWNLM0789E The IO group IO Group for virtual disk virtual disk does not have appropriate connectivity to the host(s) selected in the plan.	2007

HWNLM0790E Not considering Subsystem(s): Subsystem because not all of its back-end subsystems are configured into this IBM Spectrum Control or some are currently undetectable. This is unacceptable when an input RAID level is specified or a workload profile other than space-only is selected.	2008
HWNLM0791E The input volume you specified is already defined as a snapshot copy volume of a snapshot copy session, volume ID = volume_id, volume name = volume_name.	2008
HWNLM0792E The input volume you specified cannot be defined as a source volume in a replication session, ID = volume_id, volume name = volume_name.	2008
HWNLM0793I Replication Manager is used to manage snapshot copy sessions.	2008
HWNLM0794E The total space requirement (total space MB) cannot be met within a single storage pool for the replication session. This could be due to insufficient space in the selected storage subsystem's pool(s), or because the storage subsystem's pool(s) do not meet the requirements of the Planner input.	2009
HWNLM0796E The input volumes you specified must be within a single storage pool for the replication session.	2009
HWNLM0797E Volumes could not be created or used in the source pool of the existing snapshot copy session.	2009
HWNLM0798E The total space requirement (total space MB) cannot be met within the storage pool of the existing replication session. This could be due to insufficient space in the storage pool of the existing replication session, or because the input volumes you specified are not within the storage pool of the existing replication session.	2009
HWNLM0799E The input volumes you specified are already defined in an existing snapshot copy session.	2010
HWNLM0800W The input volume you specified is already defined as a snapshot copy volume of a snapshot copy session, volume ID = volume_id, volume name = volume_name.	2010
HWNLM0801E The input volumes you specified are not within the storage pool of the existing replication session.	2010
HWNLM0802E For a Metro Global Mirror session, you must specify at least three DS8000 subsystems with sufficient capacity.	2010
HWNLM0803E The selected input volume volume_name is already in a copy set used by this copy session.	2010
HWNLM0804E Subsystem cannot be considered. Please check both the candidate SVC and its back-end subsystems for available space and performance data for the specified time interval.	2011
HWNLM0804I The storage pool(s) provided do not satisfy the criteria for an acceptable destination pool.	2011
HWNLM0805E Host Host is a virtual machine without any Fibre Channel host port and a storage volume cannot be assigned directly to it.	2011
HWNLM0806W Host Host is a hypervisor. IBM Spectrum Control is unable to set the multipath policy on the hypervisor.	2011
HWNLM0807W Host Host is a server that is not managed by an SRA. IBM Spectrum Control is unable to set the multipath policy on the host.	2011
HWNLM0808E Ports of host Host are not connected to any fabric that is known to IBM Spectrum Control.	2012
HWNLM0809E Ports of host Host are not connected to fabrics that allow automatic zoning.	2012
HWNLM0810E The total capacity requested is {0} GB for {1} number of volumes. The individual volume size comes out to be {2} GB, which is invalid. The allowed minimum size is {3} GB and the allowed maximum size is {4} GB per volume for the given subsystem(s).	2012
HWNLM0811W No fabric information is available. All fabric-related options are ignored and no fabric configuration operation is performed.	2012
HWNLM0812E Unable to plan virtualizer provisioning task due to reaching the maximum vdisks limitation in all candidate iogroups.	2013
HWNLM0001I An integrated SAN Planner job started with schedule creator. schedule name	2013
HWNLM0002E The integrated SAN Planner job completed with errors. Message from exception: message.	2013
HWNLM0003I The integrated SAN Planner job completed.	2013
HWNLM0004W The integrated SAN Planner job completed with warnings.	2013
HWNLM0005E The integrated SAN Planner job completed with errors.	2014
HWNLM0006I Zone set zone set name created on fabric fabric wwn.	2014
HWNLM0007I Zone zone name created on fabric fabric wwn.	2014
HWNLM0008I Zone zone name added to zone set zone set name on fabric fabric wwn.	2014
HWNLM0009I A list of ports added to zone zone name for zone set zone set name on fabric fabric wwn.	2014
HWNLM0010I Activated zone set zone set name on fabric zone set name.	2015
HWNLM0011I Started to create storage volumes.	2015
HWNLM0012E The creation of storage volumes completed with errors.	2015
HWNLM0013I Completed creating storage volumes.	2015
HWNLM0014I Started to assign storage volumes to WWPNs.	2015
HWNLM0015E The assignment of storage volumes to WWPNs completed with errors.	2015
HWNLM0016I Completed assigning storage volumes to WWPNs.	2016
HWNLM0017W The command to discover volumes on host host id failed with status status.	2016
HWNLM0018W Unable to set the multipath policy on host host id due to host failure status.	2016
HWNLM0019I Completed startTransaction command on fabric fabric wwn.	2016
HWNLM0020I Completed commitTransaction command on fabric fabric wwn.	2017
HWNLM0021E The startTransaction command on fabric fabric wwn failed with return code return code.	2017
HWNLM0022E Creation for Zone set zone set name on fabric fabric wwn failed with return code return code.	2017
HWNLM0023E Creation for Zone zone name created on fabric fabric wwn failed with return code return code.	2017
HWNLM0024E Adding Zone zone name to zone set zone set name on fabric fabric wwn failed with return code return code.	2018
HWNLM0025E Adding ports to zone zone name for zone set zone set name on fabric fabric wwn failed with return code return code.	2018
HWNLM0026E Activated zone set zone set name on fabric zone set name failed with return code return code.	2018
HWNLM0027E The commitTransaction command on fabric fabric wwn failed with return code return code.	2018
HWNLM0028I Starting volume discovery on host host.	2019
HWNLM0029I Finished volume discovery on host host.	2019
HWNLM0030I Assignment(s) between Volume volume id and Host Port(s) host ports already exist, no assignment actions will happen for these paths.	2019
HWNLM0031W Since multiple Storage Resource Groups were provided as input to the plan, the newly created volumes will not be added to any Storage Resource Groups.	2019
HWNLM0032W IBM Spectrum Control is unable to set the multipath policy on host host id because it is an ESX hypervisor. After the provisioning operation completes, log in to the hypervisor or use VMware tools to set the multipath policy on the hypervisor.	2019
HWNLM0033W IBM Spectrum Control is unable to set the multipath policy on host host id because it is not managed by an SRA. After the volume or volumes are provisioning, log in to the operating system on the agentless server to set the multipath policy.	2020
HWNLM0034I Started updating agentless server configuration with disk information.	2020
HWNLM0035I Completed updating agentless server information with disk information.	2020
HWNLM0036W The volume discovery operation failed for one or more agentless servers.	2020
HWNLM0037I Disks mappings detected for volume volume by the discovery on host host on the following paths: path.	2021
HWNLM0038W No disks were detected for volume volume by the discovery on host host.	2021
HWNLM0039I After unassigning volume volume from host host, no disks were detected anymore for it by the discovery on that host.	2021

HWNLM0040W After unassigning volume volume from host host, disks mappings were still detected for it by the discovery on that host on the following paths: path.	2021
HWNLM0041I Started to copy storage volumes.	2021
HWNLM0042E The copy of storage volumes could not be completed.	2022
HWNLM0043I Completed copy of storage volumes.	2022
HWNLM0100E No Storage Subsystem(s) passed to SAN Planner.	2022
HWNLM0101E No Storage Virtualized Controller(s) passed to SAN Planner.	2022
HWNLM0102E Virtual disk(s) were selected in a plan for volumes.	2022
HWNLM0103E Volume(s) were selected in a plan for virtual disks.	2022
HWNLM0104I Not considering subsystem {0} for new storage because the user indicated so.	2023
HWNLM0105I Not considering storage pool {0} for new storage because the user indicated so.	2023
HWNLM0106I Not considering storage virtualized controller {0} for new storage because the user indicated so.	2023
HWNLM0107I Not considering managed disk group {0} for new storage because the user indicated so.	2023
HWNLM0108I Considering virtual disk: {0} of storage virtualized controller: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship (or these requirements were over-ridden by user).	2023
HWNLM0109I Considering user selected virtual disk: {0} of storage virtualized controller: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship (or these requirements were over-ridden by user).	2024
HWNLM0110I Considering user selected volume: {0} of subsystem: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship (or these requirements were over-ridden by user).	2024
HWNLM0111I The algorithm used to identify the best location for the volumes will ignore the co-location criterion.	2024
HWNLM0300I A Path Planner job started with multipath policy (multipath), redundant fabric policy (rFabric), multipath mode (mode), and (paths) number of paths.	2024
HWNLM0301I The Path Planner job completed.	2024
HWNLM0302E Multiple paths are not supported on host host name as it has only one port.	2025
HWNLM0303E No common fabrics between Host and Subsystem	2025
HWNLM0304E There is an insufficient number of possible paths between Host and Subsystem. The number of possible paths possiblePaths are less than the required number of paths paths.	2025
HWNLM0305E Cannot create redundant paths using the specified paths number of paths.	2025
HWNLM0306E There are less than two fabrics in common between host Host and storage subsystem Subsystem.	2025
HWNLM0307W No supported multipath driver was found on host host id.	2026
HWNLM0308W No corresponding multipath device was found on host host id in case of one or more volumes.	2026
HWNLM0309W Multipath policy configuration is not supported for multipath driver on host host id.	2026
HWNLM0310W Failed to set the multipath policy on host host id because for Multipath DM driver only Round Robin policy is available.	2026
HWNLM0311W The command for setting the multipath policy on host host id has failed to execute.	2026
HWNLM0312I Please check the agent log file for more details.	2027
HWNLM0313W iSCSI ports were not mapped to the storage volume because storage subsystem Subsystem does not support iSCSI connectivity.	2027
HWNLM0314W iSCSI ports were not mapped to the storage volume because storage subsystem Subsystem does not have iSCSI connectivity configured.	2027
HWNLM0315E No common connectivity exists between Host and Subsystem.	2027
HWNLM0316I Started to set the multipath policy on host host id.	2027
HWNLM0317I The multipath policy has been successfully set on host host id.	2028
HWNLM0500I The user specified maximum number of zones is user zones however the current number of zones is current zones.	2028
HWNLM0501E The maximum number of zones in a fabric policy was violated.	2028
HWNLM0502E The zone per host policy was violated.	2028
HWNLM0503E The zone per HBA policy was violated.	2028
HWNLM0504E The zone per cluster policy was violated.	2029
HWNLM0505E The zone per controller policy was violated.	2029
HWNLM0506E The zone per controller type policy was violated.	2029
HWNLM0507E The zone per fabric policy was violated.	2029
HWNLM0508I Zone Set zone set name was created.	2029
HWNLM0509I Zone zone name was created.	2029
HWNLM0510I Zone zone name was added to Zone Set zone set name.	2030
HWNLM0511I Host port port id was added to Zone zone name.	2030
HWNLM0512I Subsystem port port id was added to Zone zone name.	2030
HWNLM0513I SAN Planner started with guidance Policy (guidance), validation policies (validation), fabric WWN (fabricWWN), and using active zone set (zone set name).	2030
HWNLM0514I The Zone Planner completed.	2030
HWNLM0515W Not considering Subsystem {0} since Planner could not obtain its information from database (it could be undetectable).	2030
HWNLM0516E Could not find any subsystem with the given input. The subsystem(s) may be undetectable.	2031
HWNLM0517W Invalid candidate subsystem {0} for performance data for SAN Planner.	2031
HWNLM0518I Not considering Pool: {0} of Subsystem: {1} since it does not have enough allocatable space.	2031
HWNLM0519W Not considering Pool: {0} of Subsystem: {1} since it is not online.	2031
HWNLM0520I Not considering Pool: {0} of Subsystem: {1} because it is not Fixed Block pool.	2031
HWNLM0521I Not considering Pool: {0} of Subsystem: {1} because it is Solid State Disk pool.	2032
HWNLM0522I Not considering Pool: {0} of Subsystem: {1} because it is not a Solid State Disk pool.	2032
HWNLM0523I Not considering Pool: {0} of Subsystem: {1} because its Encryption Group is not matching with the input.	2032
HWNLM0524I Not considering Pool: {0} of Subsystem: {1} because it is not thin-provisioning enabled.	2032
HWNLM0525I Not considering Pool: {0} of Subsystem: {1} because its lock behavior is not matching with the input.	2032
HWNLM0526I Not considering Pool: {0} of Subsystem: {1} because it is not in the selected input pools list or it is already filtered out.	2033
HWNLM0527I Not considering Pool: {0} of Subsystem: {1} because its RAID level does not match with the input RAID level.	2033
HWNLM0528I Considering volume: {0} of Subsystem: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship (or these requirements were over-ridden by user).	2033
HWNLM0529W The subsystem {0} does not have performance data for the specified time interval.	2033

HWNLM0530I Not considering volume {0} for new storage because the user indicated so.	2033
HWNLM0531I Not considering the subsystem {0} because it does not support Extent Space Efficient volumes.	2034
HWNLM0532I Not considering the Subsystem: {0} because it does not support Track Space Efficient volumes.	2034
HWNLM0533I Not considering Pool: {0} of Subsystem: {1} because it does not have repository capacity available/defined for Track Space Efficient volumes.	2034
HWNLM0534I New Capacity Planning Advice Task Started....	2034
HWNLM0535I Capacity Planning Advice Task Completed.	2034
HWNLM0536I Not considering Pool: {0} of Subsystem: {1} because its backend storage RAID level does not match the input RAID level	2035
HWNLM0537I Not considering Pool: {0} of Subsystem: {1} because its backend storage is not configured into this IBM Spectrum Control or is currently undetectable and input requires a specific backend RAID level	2035
HWNLM0538I Not considering IO Group: {0} of Subsystem: {1} because it does not have two nodes associated with it	2035
HWNLM0539I Not considering IO Group: {0} of Subsystem: {1} because it has more than or equal to 2048 vdisks	2035
HWNLM0540I Not considering IO Group: {0} of Subsystem: {1} because it does not have enough available mirroring memory	2035
HWNLM0541I Not considering Volume: {0} of Subsystem: {1} because its subsystem was found to be unacceptable based on inputs	2036
HWNLM0542I Not considering Volume: {0} of Subsystem: {1} because its pool was found to be unacceptable based on inputs	2036
HWNLM0543I Not considering Volume: {0} of Subsystem: {1} because its IO group was found to be unacceptable based on inputs	2036
HWNLM0544I Not considering Volume: {0} of Subsystem: {1} because its size does not match the input	2036
HWNLM0545I Not considering Subsystem: {0} because not all of its backend subsystems are configured into this IBM Spectrum Control or some are currently undetectable. This is unacceptable when an input RAID level is specified or a workload profile other than space-only is selected	2036
HWNLM0546I Not considering Subsystem: {0} because IBM Spectrum Control does not have adequate performance data for it and/or some of its backend subsystems. For planning with workload profiles other than space-only, this subsystem and all of its backend subsystems need to have daily performance data	2037
HWNLM0547I Not considering Subsystem: {0} since none of its IO groups have enough available mirroring memory to support vdisk mirroring input	2037
HWNLM0548I Not considering Subsystem: {0} since it does not support vdisk mirroring (code level is below v4.3)	2037
HWNLM0549I Not considering Subsystem: {0} since it does not support space-efficient vdisks (code level is below v4.3)	2037
HWNLM0550I Not considering IO Group: {0} of Subsystem: {1} because it does not appropriate connectivity to hosts selected in the input	2037
HWNLM0551I Not considering Subsystem: {0} since none of its IO groups have appropriate connectivity to hosts selected in the input	2038
HWNLM0552I Not considering Volume: {0} of Subsystem {1} since it does not match input on thin-provisioning characteristics	2038
HWNLM0553I Volume size too small {0} MB. Minimum size should be {1} MB for the pool {2}. Not considering this pool.	2038
HWNLM0554I Volume size too big {0} MB. Maximum size should be {1} MB for the pool {2}. Not considering this pool.	2038
HWNLM0555I Not considering Volume: {0} of Subsystem: {1} because it is in a known replication relationship.	2039
HWNLM0556I Volume {0} can not be moved to the following candidate pool(s) due to insufficient allocatable space: {1}.	2039
HWNLM0557I The max theoretical I/O capability of storage pool {0} on subsystem {1} has not been set. Using default value {2}.	2039
HWNLM0558W The subsystem pool {0} does not have performance data for the specified time interval.	2039
HWNLM0559I Not considering IO Group: {0} of Subsystem: {1} because compression was specified, but the IO Group does not have compression active.	2039
HWNLM0560I Not considering Volume: {0} of Subsystem {1} since it is not compressed when compression was specified.	2040
HWNLM0560W Fabric agent not available to perform the zoning operation on the fabric {0}.	2040
HWNLM0561I Fabric agent available to perform the zoning operation on the fabric {0}.	2040
HWNLM0562E Fabric service exception occurred trying to check for zone control on the fabric {0}.	2040
HWNLM0563W A zone unnecessary_zone was not created, because the host already has connectivity to the storage system using existing_zone.	2040
HWNLM0564W There is already a zone named zone_name. A zone named new_zone_name was created instead.	2041
HWNLM0565W There is already a volume named vol_name. A volume named new_vol_name will be created instead.	2041
HWNLM0566I Not considering Pool: {0} of Subsystem: {1} because it is mixed pool.	2041
HWNLM0567I Not considering Subsystem: {0} for vdisk mirroring since it does not have at least two acceptable mdiskgroups from different backend subsystems	2041
HWNLM0568I Not considering I/O group {0} of subsystem {1} because it already reached the limit of 200 compressed vdisks.	2041
HWNLM0569I Not considering subsystem {0} because none of its I/O groups can be used for compressed vdisks.	2042
HWNLM0570I Not considering I/O group {0} of subsystem {1} because host definition {2} has access restrictions.	2042
HWNLM0571I Not considering storage system {0} because none of its I/O groups allow restricted access to all hosts.	2042
HWNLM801I Synchronous Refresh Configuration of the Storage Subsystem was completed Successfully for Subsystem subsystem.	2042
HWNLM802E Synchronous Refresh Configuration of the Storage Subsystem did not complete successfully for Subsystem subsystem.	2042
HWNLM803I Replication Session was created successfully session.	2043
HWNLM804E Replication Session creation failed. ReplicationManager [session].	2043
HWNLM805I Successful check for existence of session session.	2043
HWNLM806E No Such session exists with name session.	2043
HWNLM807I CopySets Added to Session successfully session.	2043
HWNLM808E CopySets addition to session failed. ReplicationManager [session].	2043
HWNLM809E CopySets creation failed. ReplicationManager [session].	2044
HWNLM810I Storage Subsystem Configuration refreshed successfully in Replication Manager ss.	2044
HWNLM811E Storage Subsystem Configuration refresh operation failed in Replication Manager [ss].	2044
HWNLM812E Fabric agent not available to perform the zoning operation on the fabric {0}.	2044
HWNLM813I Replication Session was started successfully session.	2044
HWNLM814W Replication Session start failed. ReplicationManager [session].	2045
HWNLM815W Not considering the Storage Subsystem {0} because it is not registered in TPC-R.	2045
HWNLM816I Please use Replication Manager console to schedule point-in-time copy creation.	2045
HWNLM817W No cluster partnership exists between source subsystem [srcss] and target subsystem [tgtss].	2045
HWNLM818W No path exists between source subsystem [srcss] and target subsystem [tgtss].	2045
HWNLM820I Copysets can not be added to a session when it's in Prepared/Suspend/Recover/Flash state. Please try again later after some time. State of session can be checked from Replication Manager console.	2046
HWNLM8021W IBM Spectrum Control had no information about the Fibre Channel configuration, so it did not verify fabric connectivity or change the zoning configuration.	2046
HWNOP - Storage optimizer messages	2046
HWNOP0001I The Collection phase of the integrated Storage Optimizer job jobname has started.	2047
HWNOP0002E The Collection phase of the integrated Storage Optimizer job jobname has failed with errors. Message from exception: message.	2047
HWNOP0003I The Collection phase of the integrated Storage Optimizer job jobname has completed.	2047

HWNOP0004W The Collection phase of the integrated Storage Optimizer job jobname has completed with warnings.	2047
HWNOP0005E The Collection phase of the integrated Storage Optimizer job jobname has completed with errors.	2047
HWNOP0006I Zone set zoneset created on fabric fabric .	2048
HWNOP0007I The Collection phase has started Collecion for the Analysis job jobname, on Subsystem: subsystem , Start Date: startdate , End Date: enddate	2048
HWNOP0008I The Collection phase has started Collection for the Optimization job creator, Analysis Job ID: analysisjob , on Subsystem: subsystem	2048
HWNOP0009W The subsystem subsystem does not exist in the internal database.	2048
HWNOP0010W The Performance data summary level is unknown for subsystem subsystem.	2048
HWNOP0011I Collecting Performance Data for Subsystem subsystem Storage Pool pool with summary level level for time range range	2049
HWNOP0012E Failed to bind to Performance Manager Service	2049
HWNOP0013E Database operation failed error	2049
HWNOP0014I Optimizer Service Started	2049
HWNOP0015I Optimizer Collector called for Job job with inputs input	2050
HWNOP0016I Collecting Configuration Data for Subsystem subsystem	2050
HWNOP0017I Updating Aggregated Workload Profile Data Table in Database for Subsystem subsystem	2050
HWNOP0018I The collection of configuration and performance data is finished. Now starting Analysis...	2050
HWNOP0019I The Analyzer is computing utilizations for storage subsystem subsystem.	2050
HWNOP0020I The Analyzer is updating the database with utilizations for storage subsystem subsystem.	2051
HWNOP0021I The Analyzer is reading analysis data from the database.	2051
HWNOP0022I The Analyzer queued the job on the planner.	2051
HWNOP0023I The Analyzer is done for the job.	2051
HWNOP0024I The Planner is starting the planning operation.	2051
HWNOP0025I The Planner generated number recommendations	2051
HWNOP0026I The Planner queued the job on the effector.	2052
HWNOP0027I The Effector is adding the plans in the database.	2052
HWNOP0028I The Effector is done. The Optimization job is complete.	2052
HWNOP0029I Starting Planning phase.	2052
HWNOP0030I Starting Consolidation Plan.	2052
HWNOP0031I The Storage Optimizer cannot migrate or consolidate volume volume from source pool pool on subsystem subsystem because of space or utilization threshold constraints.	2053
HWNOP0032I The Optimizer consolidated volume volume to Storage Pool pool. The Storage Pool size is size. The Max Size is maxsize.	2053
HWNOP0033I There were duration day(s) in the selected interval for subsystem subsystem. numData of these had performance data available.	2053
HWNOP0034E No data performance data is available for the subsystem subsystem and the time interval that you have selected. Please choose a new start and end date interval to analyze.	2053
HWNOP0035I volume volumes out of maxvolume volumes were not consolidated.	2053
HWNOP0036E The Optimizer job failed with the following error: error	2054
HWNOP0037I Performance Data for Subsystem subsystem Storage Pool pool will not be collected as it is of mixed pool type.	2054
HWNOP0038I The Storage Optimizer job has been queued for processing.	2054
HWNPM - Performance manager messages	2054
HWNPM0001E The specified summarization level (level) is invalid. It must be an integer value between minimum and maximum, inclusive.	2060
HWNPM0002E The specified device category (category) is invalid. It must be an integer value between minimum and maximum, inclusive.	2060
HWNPM0003E The specified device type (type) is invalid. It must be an integer value between minimum and maximum, inclusive.	2060
HWNPM0004E The specified component type (type) is invalid. It must be an integer value between minimum and maximum, inclusive.	2060
HWNPM0006E The string specified as parameter (string) exceeded its allowed length (maximum length).	2061
HWNPM0007E The value specified as parameter (value) is invalid.	2061
HWNPM0008E A required parameter is missing (null).	2061
HWNPM0010E The specified device ID (device ID) is invalid. It must conform to the pattern 'name+nameFormat'.	2061
HWNPM0011E The specified component ID (component ID) is invalid. It must be a simple WWN (16 hexadecimal characters).	2061
HWNPM0012E The specified component ID (component ID) was not found or is not unique in the IBM Spectrum Control database.	2062
HWNPM0013E The specified component ID (component ID) is invalid.	2062
HWNPM0015E Failed to retrieve the requested data because the service is unavailable.	2062
HWNPM0021E The device identifier specified as parameter (device ID) is invalid.	2062
HWNPM0090E Failed to retrieve the requested data because the service is unavailable.	2062
HWNPM0099E The requested operation failed because of an internal error.	2063
HWNPM0101E Unable to create the specified performance service instance {{0}}.	2063
HWNPM0200I This operation (operation name)on Performance Manager was successful.	2063
HWNPM0201E The device (device_id) that was passed to the method is invalid.	2063
HWNPM0202E The device category (device_category) that was passed to the method is invalid.	2064
HWNPM0203E The device type received (device_type) is invalid.	2064
HWNPM0204E The device type - HOST - that was passed to the method is not supported.	2064
HWNPM0205E The specified performance collection policy is invalid.	2064
HWNPM0209I The device type and device category are valid.	2064
HWNPM0210E Collector failed to start due to system failure.	2064
HWNPM0220E Collector failed to stop due to system failure.	2065
HWNPM0230E One or more of the specified performance collection policies are invalid.	2065
HWNPM0231W The specified performance collection policy is ignored because it conflicts with another policy in the same parameter list.	2065
HWNPM0232E The specified performance collection policy contains an unsupported interval length.	2065
HWNPM0233E The specified performance collection policy contains an unsupported frequency.	2066
HWNPM0234E The specified performance collection policy contains an unsupported duration.	2066
HWNPM0240E The attempt to update the specified performance collection policies has failed.	2066
HWNPM0241E The attempt to reset the specified performance collection policies has failed.	2066
HWNPM0242E The attempt to remove the specified performance collection policies has failed.	2067
HWNPM0249W An attempt to dynamically update one or more running performance collectors with a new performance collection policy has failed.	2067

HWNPM0250E One or more default performance collection policies are missing from the database.	2067
HWNPM0281I Performance monitoring is unavailable for resource resource_name because an agent for monitoring the resource was not defined. For IBM Spectrum Scale, the problem might occur because the data collection service cannot connect to port 9084 on the node where the collector component of the IBM Spectrum Scale performance monitoring tool is running.	2067
HWNPM0282I Performance monitoring is unavailable for resource resource_name because the associated data sources are unable to collect performance data from the resource.	2068
HWNPM0283I Performance monitoring is unavailable for resource_name because this resource does not support performance monitoring.	2068
HWNPM0284I Performance monitoring is unavailable for resource resource_name because the associated agent does not have the required level of software agent_level.	2068
HWNPM0285I Performance monitoring is unavailable for resource resource name because the associated agent is unable to fully monitor the resource.	2069
HWNPM0286I Performance monitoring is unavailable for resource resource name because the associated SMI-S provider does not have the required SMI-S support.	2069
HWNPM0287I Performance monitoring is unavailable for resource resource name because the resource or the associated agent does not support performance monitoring.	2069
HWNPM0288I Performance monitoring is unavailable for resource resource name because the resource was not probed.	2069
HWNPM0289W Performance monitoring is unavailable for resource resource_name because no agents are available.	2070
HWNPM0290E Performance monitoring is unavailable for resource resource_name because the associated agent was could not be selected.	2070
HWNPM0291I Performance monitoring is unavailable for switch resource_name because the switch has no ports.	2070
HWNPM0292I Performance monitoring is unavailable for flash resource_name because the switch was not probed using the correct agent.	2070
HWNPM0293I Performance monitoring is unavailable for FlashSystem resource_name because its SNMP agent is disabled. You can enable SNMP for FlashSystem storage systems in the FlashSystem GUI<a>.	2071
HWNPM0300E There is an exception for each device processed in a multiple devices call.	2071
HWNPM0390E A system failure occurred.	2071
HWNPM0400I This operation (operation name)on Threshold Service was successful.	2071
HWNPM0401E The device (device_id) that was passed to the method is invalid.	2072
HWNPM0410E The Performance threshold policy that was passed to the method (threshold policy)is null.	2072
HWNPM0411E The Performance threshold that was passed to the method (threshold)is null.	2072
HWNPM0412E The Performance threshold filter that was passed to the method (filter)is null.	2072
HWNPM0420E The device type received (device_type) is invalid.	2072
HWNPM0421E There is no default performance threshold policy or default threshold filter for this device.	2073
HWNPM0425E There is an exception for each device processed in a multiple devices threshold call.	2073
HWNPM0590E Performance Manager failed due to system failure.	2073
HWNPM0600E Parameter number a number of the call made to the IBM Spectrum Control Performance Manager reporting API method name of the api is invalid. The invalid value is the invalid value of parameter descriptive name of the parameter.	2073
HWNPM0601E A request to continue the data retrieval can not be performed. Information for continuing the data retrieval does not exist.	2073
HWNPM0602E Support for the device type device type name is not available in the Performance Manager reporting API function method name.	2074
HWNPM0603E The performance reporting API method_name failed at time_of_failure as a result of an internal processing exception. The Performance Manager logs contain message message_ID that describes the internal processing exception.	2074
HWNPM0604E The sort order parameter of a call to the performance reporting API method_name contains a value not included in the report columns list, at position list_item in the sort order list.	2074
HWNPM0605E Performance metadata cannot be displayed.	2074
HWNPM0606E Unable to instantiate performance reporting service service class name.	2075
HWNPM0607E An error occurred while the performance data was being retrieved.	2075
HWNPM0630E An invalid operator (operator identifier) was specified for the filter expression.	2075
HWNPM0631E An invalid first operand (operand class) was specified for the filter expression. It must be a operand class class.	2075
HWNPM0632E An invalid first operand was specified for the filter expression. The data type of the operand (data type) is invalid or unsupported.	2076
HWNPM0633E An invalid second operand (operand class) was specified for the filter expression. It must be a operand class class.	2076
HWNPM0650E The IBM Spectrum Control Performance Manager reporting API method name failed as a result of exception the related exception from a call to method method name, of the lower-level service name internal service.	2076
HWNPM0651E The configuration data needed to generate the affected volumes and hosts report for the device device name was not found in the IBM Spectrum Control database.	2076
HWNPM0652E The requested performance metrics are no longer available. Clear your browser cache to proceed.	2077
HWNPM2000I Performance monitoring is enabled.	2077
HWNPM2001E The IBM Spectrum Control Performance Manager is not operational.	2077
HWNPM2002E An initialization error occurred.	2077
HWNPM2003E Initialization of the Device server event service failed. No performance threshold exception alerts will be generated.	2077
HWNPM2004E Initialization of the product scheduler status service failed. The status of performance monitors will not be updated in the GUI.	2078
HWNPM2005E Initialization of the product configuration data service failed. Performance monitors cannot be started without this service.	2078
HWNPM2006E Initialization of the product configuration data service failed. Performance monitors cannot be started without this service.	2078
HWNPM2007E Initialization of the product counter data service failed for device type using agent type. Performance monitors will not be able to collect performance data from devices of this type.	2079
HWNPM2008E Initialization of the product metadata service failed. Performance monitors cannot be started without this service.	2079
HWNPM2009E Unable to instantiate lower level service service_class_name.	2079
HWNPM2010E Unable to instantiate the collection logic implementation service class name.	2079
HWNPM2011E Unable to instantiate the performance statistics data class class name.	2080
HWNPM2012I The product is using trace log directory log directory name.	2080
HWNPM2020W The performance monitor for device device name is not currently active, so a dynamic update of its monitor policy is not necessary.	2080
HWNPM2021W The performance monitor for device device name is not currently active, so a dynamic update of its threshold policy is not necessary.	2080
HWNPM2022E A performance monitor for device device name is already active. A new monitor for the same device cannot be started until the previous monitor completes or is cancelled.	2081
HWNPM2023W The performance monitor for device device_name is not currently active.	2081
HWNPM2024E Unable to find a monitor policy applicable to resource resource name.	2081
HWNPM2025E Unable to find a threshold policy applicable to resource resource_name.	2081

HWNPM2026I The performance monitor's primary process has failed unexpectedly. Attempting to recover from the failure.	2082
HWNPM2027I The performance monitor threshold checker has failed unexpectedly. Attempting to recover from the failure.	2082
HWNPM2028I The performance monitor purge process has failed unexpectedly. Attempting to recover from the failure.	2082
HWNPM2029I Successfully recovered from the performance monitor failure.	2082
HWNPM2030E Unable to recover from the performance monitor failure. The performance monitor for the storage resource will be shut down.	2082
HWNPM2031E The performance monitor failed due to an internal error.	2083
HWNPM2032W The performance monitor for device device name is not currently using the default monitor policy, so a dynamic update of the policy is not necessary.	2083
HWNPM2033W The performance monitor for device device name is not currently using the default threshold policy, so a dynamic update of the policy is not necessary.	2083
HWNPM2040E The device key specified for the snapshot vote (key) was not found in the database. The device does not exist.	2083
HWNPM2050E Failed to get the latest configuration data for device device_name.	2084
HWNPM2051E No performance data was collected from device device_name for the current collection interval (time_stamp) because the performance monitor was stopped.	2084
HWNPM2052E No performance data was collected from device device_name for the current collection interval due to an error. Data was last collected at time_stamp.	2084
HWNPM2053E The new performance data collected from device device_name could not be saved in the database. Increase the size of the transaction log.	2084
HWNPM2054E The new performance data collected from device device_name could not be saved in the database. Increase the size of the database lock list.	2085
HWNPM2055E The new performance data collected from resource device name could not be saved.	2085
HWNPM2056E No performance data was collected from device device name for the current performance monitor job duration. The performance monitor job status is set to 'failed'.	2085
HWNPM2057E No performance data was collected from device device_name for the current collection interval because the performance monitor was stopped.	2085
HWNPM2058E No performance data was collected from device device_name for the current collection interval due to an error.	2086
HWNPM2060W The device does not support performance management for segment pool pool ID. Only incomplete performance data can be collected for array array ID.	2086
HWNPM2061W The device does not support performance management for segment pool pool ID. Only incomplete performance data can be collected for device adapter DA ID.	2086
HWNPM2062W Invalid error message saved in database	2086
HWNPM2100E The performance monitor for resource device name cannot be started because configuration data for the resource is not available.	2087
HWNPM2101E All agents that can collect performance data for resource device name are currently non-operational.	2087
HWNPM2102E The performance monitor for resource device name cannot be started because the resource might not support the collection of performance data.	2087
HWNPM2103W Agent agent name is non-operational. Attempting to find an alternative agent.	2087
HWNPM2104I The performance monitor policy was adjusted due to agent limitations. Current values in effect are: interval-length=interval-length, frequency=frequency.	2088
HWNPM2105E The performance monitor for resource resource name failed because the resource for enabling performance data collection cannot be reached.	2088
HWNPM2106E The performance monitor for device device name failed because of errors trying to enable performance data collection on the device or device agent: error description	2088
HWNPM2107E The performance monitor for device device name failed because of unrecognized errors trying to enable performance data collection on the device or device agent: error description	2088
HWNPM2108E The performance monitor for resource resource name failed during shutdown because the resource cannot be reached for terminating data collection.	2089
HWNPM2109E The performance monitor for resource resource name failed during shutdown because of errors during termination of performance data collection: error description	2089
HWNPM2110E The performance monitor for resource resource name failed during shutdown because of unrecognized errors during termination of performance data collection: error description	2089
HWNPM2111E The performance monitor for resource resource name failed because of errors retrieving the most recent configuration data for the resource.	2090
HWNPM2112I Agent agent name was selected for performance data collection from resource resource name.	2090
HWNPM2113I The performance monitor for resource resource name is starting in an active state.	2090
HWNPM2114I The performance monitor for resource resource name is starting in a dormant state.	2090
HWNPM2115I Monitor Policy: name="policy name", creator="policy creator", description="policy description"	2091
HWNPM2116I Monitor Policy: retention period: sample data=length in days days, hourly data=length in days days, daily data=length in days days.	2091
HWNPM2117I Monitor Policy: interval length=length in seconds secs, frequency=length in seconds secs, duration=length in hours hours.	2091
HWNPM2118I Threshold Policy: name="policy name", creator="policy creator", description="policy description"	2091
HWNPM2119I Threshold Policy: retention period: exception data=length in days days.	2091
HWNPM2120I Threshold Policy: threshold name=name, component=component type, enabled=Yes or No, boundaries=critical stress boundary,warning stress boundary,warning idle boundary,critical idle boundary units.	2092
HWNPM2121I Monitor Policy: interval length=length in seconds secs, frequency=length in seconds secs, duration=continue indefinitely.	2092
HWNPM2122W No valid performance data was provided by the monitored resource. No performance data records were inserted into the database.	2092
HWNPM2123I Performance data for resource timestamp date and time was collected and processed successfully. record count performance data records were inserted into the database repository.	2092
HWNPM2124W Performance data continuity is broken. The device was possibly reset or rebooted. record count performance data records were discarded.	2093
HWNPM2125W Aggregated performance values have been computed from the remaining data records, but their accuracy cannot be guaranteed.	2093
HWNPM2126I The performance monitor for device device name is stopping because its intended duration has elapsed.	2093
HWNPM2127I The performance monitor for device device name is stopping due to a user request.	2094
HWNPM2128E The performance monitor for device device name is stopping due to an unexpected failure.	2094
HWNPM2129I The performance monitor for device device name is stopping because of a shutdown request.	2094
HWNPM2130W Failed to retrieve the latest configuration data for device device name.	2094
HWNPM2131W Performance data could not be collected for device device name, because the device or data source cannot be reached (reason reason code). The current samples are skipped.	2095
HWNPM2132W Performance data could not be collected for device device name. The current samples are skipped. (error description)	2095
HWNPM2133W Performance data could not be collected for device device name due to an unknown error. The current samples are skipped.	2095
HWNPM2134W The state of the performance monitor for resource resource name started, but the status of the performance monitor was not updated.	2096
HWNPM2135W The state of the performance monitor for device device name has changed to 'active', but could not be recorded appropriately.	2096

HWNPM2136W The performance monitor for the resource resource name generated a warning, but the status of the performance monitor was not updated.	2096
HWNPM2137W The performance monitor for the resource resource name stopped, but the status of the performance monitor was not updated.	2096
HWNPM2138W The performance monitor for the resource resource name completed the collection of data, but the status of the performance monitor was not updated.	2096
HWNPM2139W The performance monitor for the resource resource name failed, but the status of the performance monitor was not updated.	2097
HWNPM2140W The status of the performance monitor for the resource resource name was not updated.	2097
HWNPM2141E The service is unavailable because an unexpected error occurred.	2097
HWNPM2142E Performance data can't be collected for the resource resource name because the performance monitor was disabled.	2097
HWNPM2143E The performance monitor for the resource resource name was started, but the status of the performance monitor was not updated and might not be shown in the GUI.	2098
HWNPM2144W The performance data cannot be checked against the alert conditions, so no alerts can be generated.	2098
HWNPM2145I The data is being collected by the data collector: data collector host.	2098
HWNPM2146W Performance data could not be collected for device device name, the exact reason for the failure could not be determined. The current samples are skipped.	2098
HWNPM2147W Performance data could not be collected for device device name, because of a bad target (device or agent) address. The current samples are skipped.	2099
HWNPM2148W Performance data could not be collected for device device name, because of an unknown target address. The current samples are skipped.	2099
HWNPM2149W Performance data could not be collected for device device name, because of an unreachable target address. The current samples are skipped.	2099
HWNPM2150W Performance data could not be collected for device device name, because of an unresponsive target. The current samples are skipped.	2100
HWNPM2151W Performance data could not be collected for device device name, because a communication time-out for communication that uses UDP rather than TCP. The current samples are skipped.	2100
HWNPM2200I The performance monitor successfully collected the configuration data for the storage system with the following internal resources: number_of_pools pools, number_of_controllers controllers, number_of_device_adapters device adapters, number_of_ports ports, number_of_host_connections host connections, number_of_ranks ranks, number_of_arrays arrays, and number_of_volumes volumes.	2100
HWNPM2201I The performance monitor successfully collected the configuration data for the storage system with the following internal resources: number_of_io_groups I/O Groups, number_of_nodes nodes, number_of_ports ports, number_of_host_connections host connections, number_of_pools pools, number_of_managed_disks managed disks, number_of_local_disks local disks, number_of_volumes volumes, and number_of_volume_copies volume copies.	2100
HWNPM2202I The performance monitor successfully retrieved the configuration data for the switch. The following internal resources were found: number_of_trunks trunks, and number_of_ports ports.	2101
HWNPM2203I The performance monitor successfully retrieved the configuration data for the storage system. The following internal resources were found: number_of_host_connections host connections, number_of_modules modules, number_of_ports ports, number_of_pools pools, and number_of_volumes volumes.	2101
HWNPM2204I The performance monitor successfully retrieved the configuration data for the storage system. The following internal resources were found: number_of_nodes nodes, number_of_ports ports, and number_of_modules flash modules.	2102
HWNPM2205I The performance monitor successfully retrieved the configuration data for the storage system. The following internal resources were found: number_of_ports ports, number_of_controllers controllers, number_of_volumes volumes, and number_of_disks disks.	2102
HWNPM3000E There was a problem establishing the database connection.	2102
HWNPM3001E An unexpected null row was returned from a database cursor.	2102
HWNPM3002E An unexpected database exception occurred.	2103
HWNPM3003E An unexpected database exception occurred on the snapshot database tables.	2103
HWNPM3004E The snapshot ID could not be found.	2103
HWNPM3500E The current transaction has been rolled back because of a deadlock.	2103
HWNPM3501E The current transaction has been rolled back because of a timeout.	2104
HWNPM3502E The current transaction has been rolled back because the database transaction log has been exhausted.	2104
HWNPM3503E The current transaction has been rolled back because the database disk space has been exhausted.	2104
HWNPM3600E The threshold identifier parameter value : threshold ID is not valid.	2104
HWNPM3601E The target component type parameter value : component type is not valid for the threshold identifier : threshold ID passed to the affected volumes and hosts reporting function.	2104
HWNPM3602E There was a problem retrieving the performance data needed to generate the affected volumes and hosts report for the device device name.	2105
HWNPM3603E The sample volume performance data needed to generate the affected volumes and hosts report for the device device name was not found in the IBM Spectrum Control database.	2105
HWNPM3604E There are no volumes associated with the specified target component, component name, in the IBM Spectrum Control database. Therefore, the resulting Affected Volumes and Hosts report will be empty.	2105
HWNPM4000E Unable to retrieve the device agent that managed this device: device identifier.	2106
HWNPM4001E Timeout while starting performance data collection for this device: device identifier.	2106
HWNPM4002E Unable to start performance data collection for this device: device identifier.	2106
HWNPM4003E Performance data collection has already been enabled for this device: device identifier.	2106
HWNPM4004E Failed to enable performance data collection for this device: device identifier.	2107
HWNPM4005I Successfully enabled performance data collection on the storage subsystem, using device access point SMI-S provider address.	2107
HWNPM4006E An exception occurred while starting performance data collection for this device: device identifier.	2107
HWNPM4007E A timeout occurred while stopping performance data collection for this device: device identifier.	2107
HWNPM4008E Unable to stop performance data collection for this device: device identifier.	2108
HWNPM4009E Performance data collection is not enabled for this device: device identifier.	2108
HWNPM4010E Failed to disable performance data collection for this device: device identifier.	2108
HWNPM4011I Successfully disabled performance data collection on the storage subsystem, using device access point SMI-S provider address.	2108
HWNPM4012E An exception occurred while stopping performance data collection for this device: device identifier.	2109
HWNPM4013E A timeout occurred while retrieving the status of the performance data collection for this device: device identifier.	2109
HWNPM4014E Unable to retrieve the status of the performance data collection for this device: device identifier.	2109
HWNPM4015I Performance data collection is not enabled for this device: device identifier.	2109
HWNPM4016I Performance data collection is enabled for this device: device identifier.	2110
HWNPM4017E Unable to determine the status of the performance data collection for this device: device identifier.	2110
HWNPM4018E Failed to retrieve the status of the performance data collection for this device: device identifier.	2110
HWNPM4019E A timeout occurred while polling the performance statistics for this device: device identifier.	2110
HWNPM4020E Unable to retrieve the performance statistics for this device: device identifier.	2111

HWNPM4021E No performance statistics available at the current time for this device: device identifier.	2111
HWNPM4022E Failed to disable performance data collection for this device: device identifier.	2111
HWNPM4023W A set of performance statistics data was empty for this device: device identifier.	2111
HWNPM4024E An exception occurred while stopping performance data collection for this device: device identifier.	2112
HWNPM4025E Unable to retrieve storage subsystem for this device: device identifier.	2112
HWNPM4026E Failed to retrieve storage subsystem for this device: device identifier.	2112
HWNPM4027E Failed to properly initialize counter data service for this device: device identifier.	2112
HWNPM4028W Performance data cannot be collected because the security role authority of the user account user name for accessing device identifier is not sufficient.	2112
HWNPM4029W Performance data cannot be collected because the collection of performance statistics is stopped on device identifier. The security role authority of the user account user name for accessing the storage system is not sufficient to start the collection of performance statistics.	2113
HWNPM4030W Performance data cannot be collected. The performance interval device interval on device identifier is greater than the sample interval and the set security role authority of the user account user name is not sufficient to update the interval value on device identifier .	2113
HWNPM4051E Failed to obtain a reference to the Performance Manager Configuration Data Service for this device: device name.	2114
HWNPM4052E Error occurred in trying to retrieve a device agent for this device: device name.	2114
HWNPM4053E Unable to locate or retrieve the device agent that manages this device: device name.	2114
HWNPM4054E Error occurred in trying to construct the poll state information for this device: device name.	2114
HWNPM4055E Unable to construct the poll state information for this device: device name.	2115
HWNPM4056E SMI-S provider operation triggered a timeout (step timeout= step timeout value seconds, operation timeout= total timeout value seconds,).	2115
HWNPM4057E Mismatch in device identifier for this device: device name.	2115
HWNPM4058E Failed to build the parameter Map for this device: device name.	2115
HWNPM4059I Performance data collection has already been enabled for this device: device name.	2116
HWNPM4060I Performance data collection was successfully started for this device: device name.	2116
HWNPM4061E Performance data collection could not be started for this device: device name.	2116
HWNPM4062I Performance data collection successfully stopped for this device: device name.	2116
HWNPM4063W Parse exception in performance data collected this device: device name.	2116
HWNPM4064E Wrong format in performance data collected for this device: device name.	2117
HWNPM4065W number of null time stamps null time stamp(s)for performance data collected from the device were substituted by server time stamp(s).	2117
HWNPM4066W count of null operational status null Port Operational Status value(s) for performance data collected from the device was/were substituted by default value(s).	2117
HWNPM4081E A database cursor operation failed.	2117
HWNPM4082E A database connect operation failed.	2118
HWNPM4083E A database retrieve operation failed.	2118
HWNPM4084E A database operation failed.	2118
HWNPM4085E A database query operation failed.	2118
HWNPM4086W A database query gave no result rows.	2118
HWNPM4087W Missing or invalid association between SMI-S provider SMI-S provider URL and device device name. The configured SMI-S provider is inoperative, or may no longer be managing the specified device.	2119
HWNPM4091E Encountered an error during execution of a discover service process.	2119
HWNPM4092E Encountered exception during execution of a discover service process.	2119
HWNPM4093E An input business object could not be converted to a CIMInstance.	2120
HWNPM4100E Failed to initialize SVC counter data service discover service reference.	2120
HWNPM4101E Failed to initialize SVC counter data service configuration service reference.	2120
HWNPM4102E Failed to parse performance data file time stamp suffix: filename.	2120
HWNPM4103E SMI-S provider operation timeout (timeout value seconds) expired.	2121
HWNPM4104E Failed to retrieve SMI-S provider password for SVC counter data service access point: access point.	2121
HWNPM4105E Encountered an error when communicating with the device agent.	2121
HWNPM4106E Encountered invalid SVC component type: component type.	2121
HWNPM4107E Failed to create performance data object: performance data object class.	2122
HWNPM4108E TimeZone property is not defined for SVC cluster: cluster identifier.	2122
HWNPM4109E SVC cluster TimeZone property is set to unrecognized value: timezone id and name.	2122
HWNPM4110E StatisticsStatus property is not defined for SVC cluster: cluster identifier.	2122
HWNPM4111E Failed to retrieve dump filename dump from SVC node node identifier (return code = return code).	2123
HWNPM4112E IsConfigNode property is not defined for SVC node: node identifier.	2123
HWNPM4113E Caught exception while processing SVC XML performance data.	2123
HWNPM4114E SVC cluster cluster identifier has more than one configuration node.	2123
HWNPM4115E SVC cluster cluster identifier does not have a configuration node.	2124
HWNPM4116W Failed to associate SVC performance data from non-configuration node with SVC performance data from configuration node.	2124
HWNPM4117W Encountered incomplete SVC performance data sample.	2124
HWNPM4118E Firmware version information is not available for storage subsystem subsystem name. Performance data collection cannot proceed.	2124
HWNPM4119E The firmware installed on storage subsystem subsystem name (firmware version) is not supported for performance data collection. The minimum level of firmware supported for performance data collection is firmware version.	2125
HWNPM4150E Unable to retrieve storage subsystem for this device: device identifier.	2125
HWNPM4151E Unable to determine the status of any performance data collection for this device: device identifier.	2125
HWNPM4152E Performance data collection has already been enabled for this device: device identifier.	2125
HWNPM4153E Performance data collection is not enabled for this device: device identifier.	2126
HWNPM4154E Unable to start performance data collection for this device: device identifier.	2126
HWNPM4155E Failed to enable performance data collection for this device: device identifier.	2126
HWNPM4156E Unable to stop performance data collection for this device: device identifier.	2126
HWNPM4157E Failed to disable performance data collection for this device: device identifier.	2127
HWNPM4158E Unable to complete start performance data collection task for this device: device identifier.	2127
HWNPM4159E Unable to complete stop performance data collection task for this device: device identifier.	2127

HWNPM4160E Unable to complete performance data collection status query task for this device: device identifier.	2127
HWNPM4161E Performance data collection is not enabled for this device: device identifier.	2128
HWNPM4162E Unable to retrieve port performance statistics data for this device: device identifier.	2128
HWNPM4163E Unable to retrieve volume performance statistics data for this device: device identifier.	2128
HWNPM4164E Unable to retrieve rank performance statistics data for this device: device identifier.	2128
HWNPM4165E Unable to retrieve performance statistics data for this device: device identifier.	2129
HWNPM4166E Unable to complete polling for performance data collection task for this device: device identifier.	2129
HWNPM4167E Unable to retrieve a device agent for this device: device identifier.	2129
HWNPM4168E Failed attempt to use device device identifier counter data service with device different device identifier.	2129
HWNPM4169E An invalid access point of device agent URL was used to acquire the agent for this device: device identifier.	2130
HWNPM4170E The device agent's configuration for device identifier has changed from the given access point, device agent URL.	2130
HWNPM4171E Performance data collection start task timed out after time seconds for device: device identifier.	2130
HWNPM4172E Performance data collection stop task timed out after time seconds for device: device identifier.	2130
HWNPM4173E Performance data collection check status task timed out after time seconds for device: device identifier.	2131
HWNPM4174E Performance data collection poll task timed out after time seconds for device: device identifier.	2131
HWNPM4175W An error occurred while parsing statistics for port port identifier. Its statistics will be excluded.	2131
HWNPM4176W An error occurred while parsing statistics for volume volume identifier. Its statistics will be excluded.	2131
HWNPM4177W An error occurred while parsing statistics for rank rank identifier. Its statistics will be excluded.	2132
HWNPM4178E Failed to decrypt the device agent's password for device device identifier.	2132
HWNPM4179W Performance data collection is currently enabled with errors for device device identifier.	2132
HWNPM4180E Unable to retrieve key identifier value from the internal discover process.	2132
HWNPM4181W number of ports of the port statistics from the device agent were unrecognized and were not included in this sample interval.	2133
HWNPM4182W number of volumes of the volume statistics from the device agent were unrecognized and were not included in this sample interval.	2133
HWNPM4183W number of ranks of the rank statistics from the device agent were unrecognized and were not included in this sample interval.	2133
HWNPM4184E The device agent configured for this storage subsystem is not supported for this task. The current version, version number, is downlevel from from the minimum required, version number.	2133
HWNPM4185W The device agent did not return all performance statistics data for this time interval. The incomplete data is being processed.	2134
HWNPM4186W The ESS SMI-S provider did not return performance statistics data for both clusters for this time interval. The incomplete data is being processed.	2134
HWNPM4187W The device does not support performance management for pool pool ID because it contains Space Efficient Volumes. Only incomplete performance data can be collected for array array ID.	2134
HWNPM4188W The performance monitor was unable to collect performance statistics data from the device agent for the following component types: component list.	2134
HWNPM4189W number of MDisks of the MDisk statistics from the device agent were unrecognized and were not included in this data collection interval.	2135
HWNPM4190W number of nodes of the node statistics from the resource agent were unrecognized and were not included in this data collection interval.	2135
HWNPM4191W number of modules out of total number of modules module statistics could not be retrieved from the device agent due to errors, and were not included in this data collection interval.	2135
HWNPM4192W number of Drives of the drive statistics from the device agent were unrecognized and were not included in this data collection interval.	2136
HWNPM4193W number of Volume-copies of the volume-copy statistics from the device agent were unrecognized and were not included in this data collection interval.	2136
HWNPM4194W number of partitions of the partition statistics from the device agent were unrecognized and were not included in this data collection interval.	2136
HWNPM4195W number of file systems of the file system statistics from the device agent were unrecognized and were not included in this data collection interval.	2137
HWNPM4250E Failed to start the discover service for the SMI-S counter data service.	2137
HWNPM4251E Failed to start the configuration service for the SMI-S counter data service.	2137
HWNPM4252I Successfully returned access point device name for device device name.	2137
HWNPM4253I Successfully stopped SMI-S counter data service on access point access point for device device name.	2138
HWNPM4254I The SMI-S counter data service is active on access point access point for device device name.	2138
HWNPM4255I The SMI-S counter data service is inactive on access point access point for device device name.	2138
HWNPM4256I Performance statistics successfully returned on access point access point for device device name.	2138
HWNPM4257W Performance statistics not returned on access point access point for device device name.	2138
HWNPM4258E No SMI-S providers found for device device name.	2139
HWNPM4259E No storage subsystem found for device device name.	2139
HWNPM4260E Failed to initialize the polling context for device device name.	2139
HWNPM4261E Failed to retrieve the device capabilities for device device name.	2139
HWNPM4262E A database exception occurred trying to retrieve the device capabilities for device device name.	2140
HWNPM4263E A database exception occurred trying to retrieve the storage subsystem for device device name.	2140
HWNPM4264W Failed to retrieve manifest for component type.	2140
HWNPM4265E A database exception occurred trying to retrieve the Manifests for device device name.	2140
HWNPM4266E No manifests found for device device name.	2140
HWNPM4267E A database exception occurred trying to retrieve the discovery parameters for device device name.	2141
HWNPM4268E Statistics record not correctly formatted due to exception local exception string.	2141
HWNPM4269E Statistics record not correctly parsed due to exception local exception string.	2141
HWNPM4270W The block storage statistics is not formatted for device device name.	2141
HWNPM4271E The SMI-S provider found for device device name is not valid.	2142
HWNPM4272E The storage subsystem found for device device name is not valid.	2142
HWNPM4273W Discarding the stale performance statistics returned on access point access point for device device name.	2142
HWNPM4274E The SMI-S provider found for this device has changed. Please re-run SMI-S provider discovery and probe.	2142
HWNPM4300E Access to the agent or device has been denied. Ensure that valid credentials have been specified for agent agent name.	2143
HWNPM4301E The device or device agent did not respond within the allotted time (timeout value seconds).	2143
HWNPM4302E New performance data is not yet available for the device. Statistics with time stamps later than time_stamp could not be found.	2143
HWNPM4303E An agent API call (API name) failed while attempting to retrieve performance data for the device.	2143
HWNPM4304E The request for performance data could not be retrieved from the queue by the data collector.	2144
HWNPM4305W No samples were received from the data collector in the expected time. The data might still arrive automatically after connection is recovered.	2144

HWNPM4306E The data collector failed to connect to the storage management service because of invalid credentials. No performance manager data can be collected from device name until valid credentials are available.	2144
HWNPM4502E Attempt to delete a default policy.	2144
HWNPM4503E A database update operation failed.	2145
HWNPM4504E A database insert operation failed.	2145
HWNPM4505E A database delete operation failed.	2145
HWNPM4506E A database cursor operation failed.	2145
HWNPM4507E A database connect operation failed.	2145
HWNPM4508E A database retrieve operation failed.	2145
HWNPM4509E A database operation failed.	2146
HWNPM4510E A database query operation failed.	2146
HWNPM4511E A database commit operation failed.	2146
HWNPM5200E The performance manager failed to publish event even name due to exception exception.	2146
HWNPM5210E The performance manager failed to receive event from other modules.	2146
HWNPM5211E The first parameter passed to this method is null.	2147
HWNPM5212E The second parameter passed to this method is invalid.	2147
HWNPM5400E The performance data collection identifiers are not valid integers: schedule ID {0}, schedule run number {1}, job run number {2}.	2147
HWNPM5401E There was a problem establishing the database connection: {0}.	2147
HWNPM5402E There was a problem creating the new run job entry: {0}.	2147
HWNPM5403E There was a problem updating the run job entry {0}: {1}.	2148
HWNPM5404E There was a problem closing the database connection: {0}.	2148
HWNPM5405E There was a problem inserting a new run job into the database: {0}.	2148
HWNPM5406E There was a problem executing an update for run job number {0} in the database.	2148
HWNPM5407E There was a problem executing an update for run job number {0} in the database.	2149
HWNPM5408E There was a problem executing an update for run number {0} in the database.	2149
HWNPM5409I Successfully retrieved the configuration data for the elastic device. Found number of nodes Nodes and number of file systems File systems.	2149
HWNPM5410W The performance monitor could not collect performance data for the following cluster nodes: nodes names.	2149
HWNPM5411W The performance monitor could not collect performance data for the following filesystems: filesystem names.	2150
HWNPM5412E Performance statistics collection is not enabled.	1659
HWNPM5413E The process failed because the userid or password provided failed to connect to the Export Tool.	2150
HWNPM5414E The process failed because the Hitachi SVP was busy and did not return data or timed out.	2150
HWNPM5415E The process failed because the performance interval is set to 0.	2151
HWNPM5416E The process failed because the performance interval for the storage system is not supported.	2151
HWNPM5417E The process failed because the Hitachi VSP Model being monitored is not known. A Hitachi Export Tool to match it cannot be found.	2151
HWNPM5418E The process failed because the data collected is out of range.	2151
HWNPM5419E Performance data can't be collected.	2152
HWNRM - Replication manager messages	2152
HWNRM0000I Connection to Replication Manager Server successful.	2152
HWNRM0001E Communication with Replication Manager server failed.	2152
HWNRM0002E The specified port number {0} is invalid.	2153
HWNRM0003E Invalid host or port specified.	2153
HWNRM0004E Unknown host error.	2153
HWNRM0005E Failed to connect to the Replication Manager server.	2153
HWNRM0006E Status update failed in database.	2153
HWNRM0007E Unable to read status from database.	2154
HWNRM0008E Certificate file for authentication with Replication Manager not found.	2154
HWNRM0009E Replication Manager certificate error.	2154
HWNRM0010E Unable to read Replication Manager authentication certificate.	2154
HWNRM0100E The delete action failed because the session check with the Replication Manager server failed. {0}	2155
HWNRM0101E One or more storage systems cannot be deleted because they contain volumes that are defined in an active replication session. {0}	2155
HWNRM0102W The volume is in active replication session. {0}	2155
HWNRM0103W Error getting storage subsystems.	2155
HWNRM0104E Error getting storage subsystem information.	2155
HWNRM0105E Storage subsystem not found in IBM Spectrum Control database.	2156
HWNRM0106E Storage subsystem type is not supported for Replication.	2156
HWNRM0107E Volume not found in the IBM Spectrum Control database.	2156
HWNRM0108E Volume not found, volume not valid. {0}	2156
HWNRM0109E Resource not found in Replication Manager. {0}	2157
HWNRM0110E Error getting volume information.	2157
HWNRM0200E Error getting server information from the database.	2157
HWNRM0201E Error updating the server information into the database.	2157
HWNRM0011E Replication Manager server is not installed.	2158
HWNRM0012E Adding the connection to the specified storage device failed on the Replication server	2158
HWNRM0013E The connection specified to be added to the Replication server is not a valid one. Please check the parameters again	2158
HWNRM0014E A problem with the following message Exception message appeared when modifying the connection on the Replication server	2158
HWNRM0015E The connection to be modified does not exist on the Replication server	2158
HWNSS - Single sign-on User Interface messages	2158
HWNSS0001E The IBM Spectrum Control device server is down and cannot perform OS user authentication. It is still possible to perform OS user authentication against the data server, however since the device server is down the IBM Spectrum Control functionality will be limited. Among the limitations is the inability to perform SSO to other applications that rely on the presence of a lightweight third party authentication token. To proceed enter a local OS username and password.	2159
HWNSS0002E The IBM Spectrum Control device server is down and cannot perform LDAP user authentication.	2159
HWNSS0003E The single sign-on (SSO) token is missing or incorrect. Enter a valid user name and password.	2159

HWNSS0004E The web server is unavailable and cannot be used for authentication. The Device server can be used for some authentication, but IBM Spectrum Control functions will be limited. To proceed, enter either the common user name that was used to install IBM Spectrum Control or the tpcFileRegistryUser user name and password.	2159
JSS - Database messages	2160
JSS0001I Scheduler service provider started.	2160
JSS0002E Scheduler service provider initialization has failed.	2160
JSS0003I Scheduler service provider initialization successful.	2161
JSS0004I Scheduler service provider shutting down.	2161
JSS0005I Scheduler service provider shutdown complete.	2161
JSS0006E Unable to connect to repository database in class name.method name.	2161
JSS0007E SQL error preparing statement type statement for table table name in class name.method name.	2161
JSS0008E SQL error inserting into table table name in class name.method name.	2161
JSS0009E SQL error updating table table name in class name.method name.	2162
JSS0010E SQL error querying table table name in class name.method name.	2162
JSS0011E the job for computer computer name in run run number of job type job creator.job name could not be submitted due to a log table insert error.	2162
JSS0012E the job for computer computer name in run run number of job type job creator.job name could not be submitted - the delay limit was exceeded.	2162
JSS0013E the job for computer computer name in run run number of job type job creator.job name could not be submitted due to transmit error.	2163
JSS0014E run run number of job type job creator.job name could not be started due to SQL update or insert error.	2163
JSS0015E The log row for run run number of job type job creator.job name could not be updated due to an SQL error. The run failed/was successful/completed with warnings.	2163
JSS0018E The log row for the job for computer computer name in run run number of job type job creator.job name could not be updated due to an SQL error. The job completed with an exit code of exit code.	2163
JSS0019E SQL error preparing statement in class name.method name.	2164
JSS0020E Unable to find creator and name for schedule schedule ID.	2164
JSS0021E Unable to process returned job number job number The job completed with an exit code of exit code.	2164
JSS0022E Unknown request type passed to completed job handler.	2164
JSS0023E SQL error querying identifier table in class name.method name.	2165
JSS0026E Unable to schedule job type job creator.job name because fetch for computer list failed.	2165
JSS0027E The obsolete log row for job type job creator.job name could not be deleted due to an SQL error.	2165
JSS0046E the job for computer computer name in run run number of job type job creator.job name could not be started due to an agent error.	2165
JSS0051E A system error occurred.	2166
JSS0052E The agent did not respond.	2166
JSS0062W The value for the config file keyword keyword must be a number between number and number. The default (value) will be used.	2166
JSS0063E Unable to connect to repository database while attempting to delete job log rows from t_run_jobs.	2166
JSS0064E SQL error deleting job log rows from t_run_jobs.	2166
JSS0066E Unable to issue alert for job type job creator.job name because of following error.	2167
JSS0070E Agent could not be reached.	2167
JSS0071E Unable to instantiate class class name.	2167
JSS0073E the script for computer computer name in run run number of job type job creator.job name could not be submitted due to an error opening script file.	2167
JSS0074E The Scheduler service repository connection has failed and it is in auto-restart mode. It will restart automatically when the repository becomes available.	2168
JSS0075W The repository connection has been lost. The Scheduler service is terminating and will automatically restart when the repository again becomes available.	2168
JSS0076I The Scheduler service is automatically restarting after recovering its repository connection.	2168
JSS0094W The job type job named creator.name is scheduled to run once at a time in the past and will not run.	2168
JSS2003L The Storage Resource Agent version is not compatible with the server.	2168
JSS2006L The agent is disabled.	2168
JSS2007W The job for computer computer name in run run number of job type job creator.job name was not started because the Storage Resource Agent is disabled..	2169
NAD - Storage Resource Agent messages	2169
NAD0001I Connecting to hostname using protocol protocol.	1660
NAD0002W Connection to hostname failed using protocol protocol: error.	1660
NAD0003I Connected to hostname using protocol protocol.	1660
NAD0004W Possible cause: provided login information is incorrect.	2173
NAD0005E Connection to hostname failed using protocol protocol: error message.	1660
NAD0006E Exception thrown for method method name: error message.	1661
NAD0007I Closing connection to hostname.	1661
NAD0008E Invalid protocol protocol passed to method name.	1661
NAD0009E Cannot connect to host name. Remote host is running in a non-global application container.	2173
NAD0010E Invalid parameter(s) parameter name passed to method name.	1661
NAD0011I Validating GUID on remote machine: host name.	2174
NAD0012I GUID value validated on remote machine: host name.	2174
NAD0013I Installing GUID on remote machine: hostname.	1661
NAD0014I GUID successfully installed on remote machine: hostname.	1661
NAD0015I GUID not found on remote machine: host name.	2174
NAD0016E Could not copy GUID package to remote host: host name.	2174
NAD0017E The GUID on remote machine: host_name could not be validated.	2174
NAD0018E Command on remote machine: host name failed. Error code = value executing command value.	1661
NAD0019E Parameter parameter passed to method is null or 0 length.	1661
NAD0020I Host validation succeeded.	2175
NAD0021E Host validation failed.	2175
NAD0022E Cannot contact remote host due to invalid credentials, check logs for additional information. Host validation failed.	2175
NAD0023E There is not enough space on the remote machine. Host validation failed.	2175

NAD0024E It was not possible to determine the available space on the remote machine. Host validation failed.	2176
NAD0025E The specified directory could not be accessed. Host validation failed.	2176
NAD0036E Failed to copy package name to remote machine.	2176
NAD0037E Cannot cleanup remote machine directory: directory path.	2176
NAD0038E Failed to install agent on host name.	2176
NAD0039I Agent successfully installed at install location.	2176
NAD0040E Agent install exited with an error code: value.	2177
NAD0041E Failed to get agent bundle locations.	2177
NAD0042E Exception occurred while retrieving IPs for Data Server machine : host name.	2177
NAD0043I Installing agent at install location.	2177
NAD0044I Agent registration to Data Server completed successfully.	2177
NAD0045I Validation of host name has started.	2177
NAD0046E GUID cannot be null value for registration of agent.	2177
NAD0047E Exception while reading authentication information.	2178
NAD0048E Probe did not start successfully on agent host address, error code returned = value.	2178
NAD0049I Running probe on agent host address.	2178
NAD0050E Exception occurred while running probe on agent host name : exception message.	2178
NAD0051I Successfully started probe on agent host address.	2178
NAD0052E Error in receiving remote file name from host address.	2178
NAD0053E Exception occurred while receiving remote file name from host address: exception message.	2179
NAD0054E The directory install_location is not empty on the remote machine. Host validation failed.	2179
NAD0055E Failed to connect to remote host host.	1661
NAD0056E Error uninstalling agent at host address:install location.	2179
NAD0057E Error uninstalling agent at host address:install location error message.	2179
NAD0058I Performing agent upgrade on remote machine host address.	2180
NAD0059I Upgrade on host address succeeded.	2180
NAD0059E Upgrade on host address failed: error message.	2180
NAD0070I Updating Langpacks on remote machine host address.	2180
NAD0071W Requested Langpack language pack name not found on local machine.	2180
NAD0072I Langpacks updated on remote machine host address.	2180
NAD0073E Error updating LangPacks on remote machine host name : error message.	2180
NAD0074W No langpacks found on local machine, update not performed.	2181
NAD0075E The Agent is already installed on host host_name. Host validation failed.	2181
NAD0076E Failed to configure Auto-upgrade feature.	2181
NAD0077E Schedule file name file cannot be created because the schedule file is a directory.	2181
NAD0078E Cannot create file name.	2181
NAD0079E Cannot delete file name.	2182
NAD0080W Auto-upgrade feature was already enabled.	2182
NAD0081W Auto-upgrade feature was already disabled.	2182
NAD0082E Exception has been encountered: exception trace.	2182
NAD0083W Following Storage Resource Agent information is invalid agent information.	2182
NAD0084W Host location has no match from IBM Spectrum Control cached non-daemon based Storage Resource Agent list. Agent might not be non-Daemon based Storage Resource Agent.	2183
NAD0085I Current non-Daemon based Storage Resource Agents in IBM Spectrum Control cache are agents list .	2183
NAD0086E Unable to connect to Windows Domain Controller host name .	2183
NAD0087E Unable to determine source package home directory.	2183
NAD0088E Error deploying package zip file.Error message.	2183
NAD0089E Error deleting binary executable file.Error message.	2183
NAD0090E No connection exists to remote machine.	2184
NAD0091E Error executing binary executable file.Error message.	2184
NAD0092E Unable to retrieve remote temporary directory location.	2184
NAD0093E Unable to retrieve domains from host name.	2184
NAD0094E Unable to retrieve host list for domain name.	2184
NAD0095E The platform OS_name is not supported for Storage Resource Agents.	2185
NAD0096E Authentication failed due to invalid credentials or insufficient access privileges.	2185
NAD0097I Opening connection to hostname.	1662
NAD0098I Copying agent files on remote machine...	2185
NAD0099I Installing agent on host host name, in install location directory with force option.	2185
NAD0100E Agent command did not run successfully on agent host name, error code returned = value.	2185
NAD0101I Running agent command on agent host name.	2186
NAD0102E Exception occurred while running command on agent host name : error message.	2186
NAD0103E Error in deleting file remote data file from host host address.	2186
NAD0104E Exception occurred while deleting file remote data file from host host address: exception trace.	2186
NAD0105E Failed to lock Agent file.	2186
NAD0106E Failed to stop Probe.	2186
NAD0107E Failed to stop Agent.	2187
NAD0108E Failed to create registry entry on Agent machine.	2187
NAD0109E Failed to extract file.	2187
NAD0110E Failed to create configuration.	2187
NAD0111E Failed to stop Agent service.	2187
NAD0112E Agent service already exist.	2187
NAD0113E Failed to create Agent service.	2188

NAD0114E Failed to start Agent service.	2188
NAD0115E Failed to start Probe.	2188
NAD0116E Failed to create Agent service.	2188
NAD0117E File extraction needs more space.	2188
NAD0118E Failed to open archive file.	2189
NAD0119E Missing upgrade files.	2189
NAD0120E Failed to extract files.	2189
NAD0121E Failed to stop probe in upgrade process.	2189
NAD0122E Failed to stop Agent in upgrade process.	2189
NAD0123E Failed to start Agent in upgrade process.	2189
NAD0124E Failed to register Agent to Server.	2190
NAD0125E Extraction needs more space.	2190
NAD0126E Failed to open archive file.	2190
NAD0127I Uninstalling agent from host host address at location path.	2190
NAD0128E Failed to start process.	2190
NAD0129E Error in opening catalog file.	2190
NAD0130E Probe failed.	2191
NAD0131E The installation process could not write files to the directory install_location. Host validation failed.	2191
NAD0135E The certificate file file_name was not found on host_name.	2191
NAD0136E The port number of the agent is invalid.	2191
NAD0137E The port number port_number is in use on the remote machine.	2191
NAD0138E Invalid character "invalid_character" found in install_location.	2192
NAD0139E The User userID does not have sufficient administrator privileges.	2192
NAD0140I Cannot ping to host host name.	2192
NAD0141E Either the remote host could not be contacted due to invalid credentials or the machine is not reachable. Check logs for more information. Host validation failed.	2192
NAD0142E Deployment of Agent failed, error creating startup scripts.	2193
NAD0143E Cannot change agent from Daemon based type to non-Daemon based type with force installation, host validation failed.	2193
NAD0144E Agent type can not be changed between non-daemon-based and daemon-based with force installation. Host validation failed.	2193
NAD0145E Cannot get version information from agent on host .	1663
NAD0146E The connection to remote machine failed because the Remote Execution and Access component was unable to create a temporary directory on the remote machine. Remove unneeded ~CSRI* directories in the remote machine's temporary directory.	1663
NAD0147E The daemon-based agent on host_address could not be stopped.	2194
NAD0148E Daemon based agent on host address failed to start.	2194
NAD0149E Runtime files of agent host address are missing or corrupted.	2194
NAD0150E Agent is defected.	2194
NAD0151E The original install location path cannot be changed.	2194
NAD0152E Initialization of TCP/IP failed while creating socket.	2195
NAD0153E Agent service could not be stopped at uninstallation time.	2195
NAD0154E Invalid server name passed at uninstallation time.	2195
NAD0155E Port passed at installation time for agent service is in use.	2195
NAD0156E The server host_address cannot be reached because the host name or IP address is not recognized.	1663
NAD0157E The server host_name cannot be contacted. The server might be down, unreachable due to network problems, or the SSH credentials might be invalid.	1663
NAD0158E GUID value of machine host_name duplicates GUID on agent machine host_address.	2196
NAD0160E Agent is already registered with Server on machine host name location path.	2196
NAD0161E The install location is not an absolute path.	2196
NAD0162E Services script did not run on the host_address server : exception_message.	2196
NAD0163I Services files collected from host address machine can be found in local service file archive.	2197
NAD0164E Cannot collect service data from the host_address server.	2197
NAD0165E Cannot copy the services_file service archive file from the host_address server.	2197
NAD0166E Cannot create the local_service directory for storing service archive file.	2197
NAD0167I Running scan on agent host address.	2197
NAD0168E Exception occurred while running scan on agent host address : exception message.	2197
NAD0169I Successfully started scan on agent host address.	2198
NAD0170E Scan did not start successfully on agent host address, error code returned = value.	2198
NAD0171E Failed to copy remote file name file to host name machine.	2198
NAD0172E Failed to extract remote file name file.	2198
NAD0173I Running NetAppImportQuota on agent host address.	2198
NAD0174E Exception occurred while running NetAppImportQuota on agent host address : error message.	2198
NAD0175I Successfully started NetAppImportQuota on agent host address.	2199
NAD0176E NetAppImportQuota did not start successfully on agent host address, error code returned = value.	2199
NAD0180I Installing re-distributable package on .	1662
NAD0181I Install of re-distributable package on succeeded.	1662
NAD0182E Failed to install re-distributable package on .	1662
NAD0183I Validating re-distributable package on host name.	2199
NAD0184I Validation of re-distributable package succeeded.	2199
NAD0185E Validation of re-distributable package failed.	2200
NAD0186I Trying to locate package TIVguid using pkginfo ...	1662
NAD0187I Package TIVguid is not installed.	1662
NAD0188I Checking TIVguid default install path : path ...	1662
NAD0189E Command is not valid.	2200
NAD0190E Provided option is not valid.	2200

NAD0191E Arguments are not valid.	2200
NAD0192E Value for one of the arguments is missing.	2201
NAD0193E Localized string missing in message file.	2201
NAD0194E Probe process is already running.	2201
NAD0195E Failed to open file for writing.	2201
NAD0196E Failed to close file.	2201
NAD0197E Logfile was not specified.	2201
NAD0198E Tracing failed.	2202
NAD0199E Cannot start a new Probe process because another one is already running.	2202
NAD0200I Validating user on remote machine host name.	2202
NAD0201I Validation of user user succeeded.	2202
NAD0202E The validation of user credentials for user failed.	2202
NAD0203E No tracing.	2202
NAD0204E An internal error occurred in the agent process.	2203
NAD0205E The provided socket is invalid.	2203
NAD0206E Failed to start the service.	2203
NAD0207E The file was not found.	2203
NAD0208E Not enough memory to run agent.	2203
NAD0209E The agent process did not start after the upgrade was finished.	2204
NAD0210E The installation directory was not valid when trying to upgrade.	2204
NAD0211E The Probe is currently busy.	2204
NAD0212E The data file was not found.	2204
NAD0213E The exit code is not in the output file.	2204
NAD0214E Failed to send the job status.	2204
NAD0215E Failed to copy the certificate files.	2205
NAD0216E Failed to create directory.	2205
NAD0217E Failed to remove directory.	2205
NAD0218E Failed to execute the command.	2205
NAD0219E Failed to convert wide characters.	2205
NAD0220E The installation directory is not valid.	2206
NAD0221E No server name was provided.	2206
NAD0222E An error occurred while trying to remove entries from the configuration file.	2206
NAD0223E Failed to stop Probe process at uninstallation time.	2206
NAD0224E Failed to remove registry entry at uninstallation time.	2206
NAD0225E Failed to remove service entry at uninstallation time.	2206
NAD0226E The agent files are corrupted.	2207
NAD0227E Failed to install the GUID.	2207
NAD0228E A storage resource agent with a different runtime operation mode, daemon / non-daemon, is already installed at the specified location.	2207
NAD0229E Not enough disk space available for installation.	2207
NAD0230E An installation is already in progress.	2207
NAD0231E Cannot get the server name from the server.	2208
NAD0232E The installation directory is not empty.	2208
NAD0233E The parameter server name is missing.	2208
NAD0234E The parameter server port is missing.	2208
NAD0235E The parameter server ip is missing.	2208
NAD0236E The parameter agent port is missing.	2208
NAD0237E The parameter installation directory is missing.	2209
NAD0238E Failed to send the Probe results.	2209
NAD0239E Failed to initialize the agent.	2209
NAD0240E The service port number is missing.	2209
NAD0241E Get data file statistics failed.	2209
NAD0242E Get data file read failed.	2209
NAD0243E Failed to send the data to the server.	2210
NAD0244E Failed to receive the data from the server.	2210
NAD0245E The full path was not specified. (copy file function)	2210
NAD0246E Create file failed. (copy file function)	2210
NAD0247E Write file failed. (copy file function)	2210
NAD0248E Open file failed. (copy file function)	2210
NAD0249E Read file failed. (copy file function)	2211
NAD0250E The UCS conversion failed.	2211
NAD0251E The server connection failed.	2211
NAD0252E Please check the OS level.	2211
NAD0253W No error message defined for error code: value.	2211
NAD0254E Registry entry not found.	2212
NAD0255E Insufficient space to copy file file name on remote machine host name in location location.	2212
NAD0256E Cannot get the available space on remote machine host name : error message.	2212
NAD0257E OS configuration error encountered. Please do a local install with debug set to MAX.	2212
NAD0258E Port number is in use on remote machine. Stop process manually or select another port number.	2212
NAD0259W Unable to determine Storage Resource Agent version on host . Fabric Discovery will not be invoked.	1662
NAD0260I Agent is active.	1664
NAD0261I Agent shutdown successfully.	2213
NAD0262I Successfully started agent	2213

NAD0263W The probe of the agent on host name did not complete in the allocated amount of time.	2213
NAD0264I The agent on host name is being probed.	2213
NAD0265I After the probe has completed, the new volumes on the host will display in the IBM Spectrum Control GUI.	2213
NAD0266I The probe of the agent on host name completed.	2214
NAD0267W The probe of the agent on host name completed with return code value.	2214
NAD0268I To display newly assigned volumes in the IBM Spectrum Control GUI, start a new probe after the currently running probe completes. Make sure the new probe completes without errors.	2214
NAD0269I To display newly assigned volumes in the IBM Spectrum Control GUI, you must start a new probe. Make sure the new probe completes without errors.	2214
NAD0270W The Storage Resource Agent on host name is disabled and will not process any requests.	2214
NAD0271W The connection to the Storage Resource Agent on host name was not established because the agent is disabled.	2215
NAD0272W The connection to the Storage Resource Agent on host name was not established. Retrying using the IP address.	1664
NAD0273E The connection to remote machine failed because the Remote Execution and Access component was unable to create a temporary directory on the remote machine. Remove unneeded ~CSRI* directories in the remote machine's temporary directory.	2215
NAD0274E An SSH certificate certificate name already exist.	1664
NAD0275E Failed to connect to remote host hostname and port. Failed to establish a secure connection.	1664
NAD0276E Failed to connect to remote host hostname and port. Failed to establish a secure connection because the SSL handshake failed.	1664
NAD0277E Failed to connect to remote host hostname and port. Failed to establish a secure connection because of an invalid SSL key.	1664
NAD0278E Failed to connect to remote host hostname and port. Failed to establish a secure connection because the identity of the peer could not be verified.	1665
NAD0279E Failed to connect to remote host hostname and port. Failed to establish a secure connection because of an SSL protocol error.	1665
NAD0280E The installation failed for the Microsoft VC++ Redistributable package on host. The return code is value.	2217
NAD0281E The Storage Resource agent cannot be deployed because of insufficient space or other issues on the target system. The error is: error message.	1665
NAG - Storage Agent Resource messages	2217
NAG0008E Please enter a user with administrative privileges for the tree.	2217
NAG0009E The user password is required.	2217
NAG0030E Please enter a user with administrative privileges on the filer.	2217
NAG0108E The Server Name is required.	2218
NAG0119E No NDS trees have been located in your installation or they have not yet been assigned a login and password.	2218
NAS - Storage Agent Resource messages	2218
NAS0003E Unable to contact the agent on host host address.	2218
NAS0004E NetApp quotas: Error gathering list of agents.	2219
NAS0005I No filers.	2219
NAS0006W No agents available to retrieve quotas from filer.	2219
NAS0008I Gathering quotas through agent host address.	2219
NAS0009E No agents available.	2219
NAS0010E Bad requestData.	2219
NAS0011E Not all agents reported back within allotted time.	2219
NAS0012W Filer filer : No quotas retrieved.	2220
NAS0013E DB error saving/checking quotas.	2220
NAS0014E *** INTERNAL ERROR ***.	2220
NAS0015I Filer filer : Quota defined for volume volume but we have no record of that volume.	2220
NAS0016W Quota quota creator.quota name has no consumers.	2220
NAS0017E DB error saving/checking quotas (filer = filer).	2220
NAS0018I NetApp Quota job completed successfully.	2221
NAS0019I NetApp Quota job completed with WARNINGS.	2221
NAS0020I NetApp Quota job completed with ERRORS.	2221
NAS0021I Processed quotas for filer filer.	2221
NAS0022E NetApp Quota is still associated with a schedule, unable to delete.	2221
NAS0023E NAS server server name is already manually registered to a windows domain.	2221
NAS0024E Network server server name is already manually registered.	2221
NAS0025E An agent is already installed on server name. It cannot be manually entered.	2222
NAS0028E Network name of filer could not be determined.	2222
SAA - Storage Resource Agent - Storage Subsystem messages	2222
SAA0001E SYMAPI error error code -- error text.	2222
SAA0002E Error connecting to SYMAPI database (mode = mode).	2223
SAA0003E Error synching Symmetrix ID.	2223
SAA0004E Symmetrix Symmetrix ID: Unexpected volume name: volume name.	2223
SAA0005E SymDevList(Symmetrix ID) failed.	2223
SAA0006E SymDiskList(Symmetrix ID) failed.	2223
SAA0007W Symmetrix Symmetrix ID: SymDiskShow(Symmetrix volume) failed.	2224
SAA0008I (Disk disk , Hyper hyper).	2224
SAA0009W No parity hyper found for RAID group 0xRAID group.	2224
SAA0010W Symmetrix Symmetrix ID: SymDevShow(Symmetrix volume) failed.	2224
SAA0011I (Meta-component number of volume volume).	2224
SAA0012W Symmetrix Symmetrix ID: Volume <Symmetrix volume> not found or already used.	2224
SAA0013W Symmetrix Symmetrix ID: Hyper not found, volume remote sequence number (instance, bus number, target, partition).	2224
SAA0014W Symmetrix Symmetrix ID: SymDevShow(Symmetrix number) failed.	2224
SAA0015W Symmetrix Symmetrix ID: Volume Symmetrix volume does not contain hyper (device name, device number, disk ID, hyper number).	2225
SAA0016W Symmetrix Symmetrix ID: Volume Symmetrix name is actually a meta-component of Symmetrix volume.	2225
SAA0017W Symmetrix Symmetrix ID: Volume Symmetrix name has no hypers.	2225
SAA0018I SYMAPI version: ID.	2225
SAA0019E SymShow(Symmetrix ID) failed.	2225
SAA0020E SymDiscover failed.	2225

SAA0021E SymList failed.	2226
SAA0022I Storage Subsystem subsystem name (subsystem alias) will be probed.	2226
SAA0023W Unsupported storage subsystem type: subsystem type.	2226
SAA0024E Failed to load SYMAPI.	2226
SAA0025W Symmetrix ID: Remote Symmetrix remote ID: Unexpected volume name: name.	2226
SAA0026W Symmetrix ID: Volume number (volume number) exceeds array capacity (length).	2226
VPLG - VASA provider messages	2227
VPLUG0001E The connection was refused because the IBM Spectrum Control server is not available at the specified host and port.	2227
VPLUG0002E Unable to find the IBM Spectrum Control server configuration information.	2227
VPLUG0003E Unable to access the IBM Spectrum Control server configuration information.	2228
VPLUG0004E The service class service_class_name does not exist.	2228
VPLUG0005E An error was received from the IBM Spectrum Control server: Key:message_key; Message:message_text; Status code: status_code	2228
VPLUG0006E An error was received from the IBM Spectrum Control server: message	2228
VPLUG0007E An error was encountered while retrieving the list of capacity pools.	2228
VPLUG0008E Authentication to the IBM Spectrum Control server failed.	2229
VPLUG0009E The specified host is not known.	2229
VPLUG0010E The specified host cannot establish a secure connection.	2229
VPLUG0011E An error was encountered when retrieving the list of service classes.	2229
VPLUG0012E An error was encountered while provisioning block storage.	2229
VPLUG0013E An error was encountered while provisioning file storage.	2230
VPLUG0014E The provisioning task timed out.	2230
VPLUG0015E An error was encountered while retrieving the status of the provisioning task.	2230
VPLUG0016E An error was encountered while retrieving the list of WWPNS.	2230
VPLUG0017E An error was encountered when retrieving the file storage NFS options.	2230
VPLUG0018E An error was encountered while checking administrator privilege.	2231
VPLUG0019E An error was encountered while retrieving fabric port information.	2231
VPLUG0020E An error was encountered while retrieving the IBM Spectrum Control port.	2231
VPLUG0021E An error occurred while retrieving virtual disk information from IBM Spectrum Control.	2231
VPLUG0022E An error was encountered while retrieving volume performance information from IBM Spectrum Control.	2231
VPLUG0023E An error was encountered while retrieving storage devices list from vCenter.	2232
VPLUG0024E An error occurred while retrieving virtual machine virtual disk information from vCenter.	2232
VPLUG0025E An error was encountered while retrieving available datastore names from vCenter.	2232
VPLUG0026E An error was encountered while retrieving the Lun WWN.	2232
VPLUG0027E An error was received from the IBM Spectrum Control server.	2232
VPLUG0028E An error was received from the IBM Spectrum Control server: Key:message_key; Message:message_text	2233
VPLUG0029E An error was encountered while retrieving roles for current user.	2233
VPLUG0030E The user {0} does not have the minimum role (Administrator, Monitor or External application) needed to access IBM Spectrum Control.	2233
VPLUG0031E Registration of the IBM Spectrum Control VASA provider could not be completed.	2233
VPLUG0032E The IBM Spectrum Control VASA provider is already registered for server server_name. Register IBM Spectrum Control VASA provider manually to update credentials.	2233
VPLUG0033E One or more third party VASA provider(s) are already registered with the vCenter. IBM Spectrum Control VASA provider was not registered. Register IBM Spectrum Control VASA provider manually.	2234
VPLUG0034E Automatic registration of IBM Spectrum Control VASA provider is not supported for vCenter server version 5.0 and previous.	2234
VPLUG0035E Error encountered while saving IBM Spectrum Control server configuration. Incorrect value for user name or password.	2234
VPLUG0036E Error encountered while saving IBM Spectrum Control server configuration. Invalid host name message.	2234
VPLUG0037E Error encountered while saving IBM Spectrum Control server configuration. Invalid port message.	2234
VPLUG0038E Error encountered while loading IBM Spectrum Control server configuration. Incorrect value for user name or password.	2235
VPLUG0039E Error encountered while loading IBM Spectrum Control server configuration. Invalid port message.	2235
VPLUG0040E An error was encountered while the server configuration information for IBM Spectrum Control was saved. The current session is invalid.	2235
VPLUG0041E An error was encountered while checking administrator privilege. The current session is invalid.	2235
Publications	2235
Accessing publications online	2236
IBM Redbooks	2236
Translation	2236
Providing feedback about publications	2236
Legal notices	2237
Trademarks	2238
Privacy policy considerations	2238
Glossary	2238
A	2238
B	2239
C	2239
D	2240
E	2240
F	2241
G	2242
H	2242
I	2243
J	2243
K	2243
L	2243
M	2243

N	2244
O	2244
P	2244
R	2245
S	2245
T	2247
U	2247
V	2247
W	2248
Z	2248

IBM Spectrum Control documentation for V5.4.9

IBM Spectrum® Control centralizes, automates, and simplifies the management of complex and heterogeneous storage environments. Use the IBM Spectrum Control documentation to help you with your storage management tasks.

What's new for IBM Spectrum Control 5.4.9

Learn about the latest changes that are available in IBM Spectrum® Control 5.4.9.

Product fixes and security updates

In every release of IBM Spectrum Control, IBM fixes issues that are reported by our internal teams and by our customers. Along with fixes, improving the operation and security of the product is also a priority. See what's been updated and fixed in this release:

- [Fixes that are included in this release](#)
- [Security bulletins for IBM Spectrum Control](#)

Limitations and known issues

For information about limitations and known issues that might occur when you use IBM Spectrum Control 5.4.9, see [Limitations and known issues for IBM Spectrum Control](#).

Related information

- [Change History for IBM Spectrum Control](#)
- [Sponsor user program](#)
- [Beta program](#)
- [Suggest improvements to the product](#)
- [Change History for IBM Spectrum Control](#)
View a history of the changes in previous versions of IBM Spectrum Control.
- [Discontinued features in IBM Spectrum Control](#)
View a list of the features that were discontinued in different releases of IBM Spectrum Control.
- [Sponsor user program](#)
Sponsor users interact directly with designers and developers to improve the user experience and to help shape the future of the overall storage portfolio.
- [Beta program](#)
The IBM Spectrum Control Beta is a continuous program. It gives you a first look at upcoming features, a chance to influence design, an opportunity to test the new features in your own environment, and a direct voice into the product development process.
- [Collaborating with the team](#)
Collaborate with the IBM Spectrum Control team to help improve the product.

Change History for IBM Spectrum Control

View a history of the changes in previous versions of IBM Spectrum® Control.

- [Changes in 5.4.9](#)
- [Changes in 5.4.8](#)
- [Changes in 5.4.7](#)
- [Changes in 5.4.6](#)
- [Changes in 5.4.5](#)
- [Changes in 5.4.4](#)
- [Changes in 5.4.3](#)
- [Changes in 5.4.2](#)
- [Changes in 5.4.1](#)
- [Changes in 5.4.0](#)
- [Changes in 5.3.7](#)
- [Changes in 5.3.6](#)
- [Changes in 5.3.5](#)
- [Changes in 5.3.4](#)
- [Changes in 5.3.3](#)
- [Changes in 5.3.2](#)
- [Changes in 5.3.1](#)
- [Changes in 5.3.0](#)
- [Changes in 5.2.17](#)
- [Changes in 5.2.16](#)
- [Changes in 5.2.15](#)
- [Changes in 5.2.14](#)
- [Changes in 5.2.13](#)
- [Changes in 5.2.12](#)

- [Changes in 5.2.11](#)
- [Changes in 5.2.10](#)
- [Changes in 5.2.9](#)
- [Changes in 5.2.8](#)

Changes in 5.4.9

The following updates are available in IBM Spectrum Control 5.4.9:

Product fixes and security updates

- [Fixes that are included in 5.4.9](#)
- [Security bulletins for IBM Spectrum Control](#)
- [Limitations and known issues](#)

Changes in 5.4.8

The following updates are available in IBM Spectrum Control 5.4.8:

More security when you want it

Security is an important part of protecting an organization's assets and data. To provide another level of security when you want it, you can now enter a password for the SSH User that IBM Spectrum Control uses to connect with IBM Spectrum Virtualize storage systems that run firmware levels 8.5.0.0 or later.

Product fixes and security updates

See the fixes and updates in this release:

- [Fixes that are included in 5.4.8](#)
- [Security bulletins for IBM Spectrum Control](#)
- [Limitations and known issues](#)

Changes in 5.4.7

The following updates are available in IBM Spectrum Control 5.4.7:

Product fixes and security updates

See the fixes and updates in this release:

- [Fixes that are included in 5.4.7](#)
- [Security bulletins for IBM Spectrum Control](#)
- [Limitations and known issues](#)

Db2® 11.1 is no longer supported

Db2 11.1 is out of service and no longer supported by IBM Spectrum Control:

- For information about how to upgrade Db2 to Db2 11.5 or later, see [Example of upgrading Db2 in a Windows environment](#).
- For a complete list of supported Db2 versions, go to <http://www.ibm.com/support/pages/node/6249361#DB>.

Changes in 5.4.6

The following updates are available in IBM Spectrum Control 5.4.6:

Product fixes and security updates

See the fixes and updates in this release:

- [Fixes that are included in 5.4.6](#)
- [Security bulletins for IBM Spectrum Control](#)
- [Limitations and known issues](#)

Changes in 5.4.5

The following updates are available in IBM Spectrum Control 5.4.5:

Enhanced OS support

You can now install IBM Spectrum Control and Storage Resource agents on servers that run AIX® 7.3 (64 bit). This enhanced support is available in IBM Spectrum Control 5.4.5.2 or later. For more information, see the following pages:

- [Release information for 5.4.5](#)
- [Installing IBM Spectrum Control in a single-server AIX or Linux® environment](#)
- [OS support for IBM Spectrum Control](#)

Product fixes and security updates

See the fixes and updates in this release:

- [Fixes that are included in 5.4.5](#)
- [Security bulletins for IBM Spectrum Control](#)
- [Limitations and known issues](#)

Changes in 5.4.4

The following feature was available in IBM Spectrum Control 5.4.4:

A new subscription option is available for IBM Virtual Storage Center

New customers can now subscribe to IBM Virtual Storage Center. With a subscription, you can pay for your license on a monthly basis with a flexible term length and an initial lower cost. You also receive IBM's Software Subscription and Support. To sign up and manage the subscription contract, you can use the familiar Passport Advantage® (PA) system. For more information about subscription licenses with IBM, see <https://www.ibm.com/software/passportadvantage/subscriptionlicenses.html>.

[View more information about the changes in 5.4.4](#)

Changes in 5.4.3

The following feature was available in IBM Spectrum Control 5.4.3:

IBM Spectrum Control now comes with IBM Spectrum Protect™ Snapshot

Protecting the application and data that drives your business and is critical to your operations is important to us. That's why IBM Spectrum Control includes a license for IBM Spectrum Protect Snapshot, which provides fast, application-aware backups and restores, leveraging advanced snapshot technologies to help improve your cyber resiliency.

IBM Spectrum Protect Snapshot is available with the following editions of IBM Spectrum Control:

- IBM Spectrum Control and IBM Spectrum Control Select Edition
- IBM Virtual Storage Center and IBM Virtual Storage Center Entry

[View more information about the changes in 5.4.3](#)

Changes in 5.4.2

The following feature was available in IBM Spectrum Control 5.4.2

Update for IBM Virtual Storage Center

The edition "IBM Virtual Storage Center Entry Edition" was renamed to "IBM Virtual Storage Center Starter Edition"

Along with the renaming, capacity limits in IBM Virtual Storage Center Starter Edition were increased to better meet your needs. Now, you can use the IBM Virtual Storage Center Starter Edition to manage up to 1,500 TiBs of usable capacity. That's an increase of 1,000 TiBs over the previous limit! Speaking of limits, the previous limit of 4 I/O groups was removed, so you're no longer restricted by the number of I/O groups in your environment.

[View more information about the changes in 5.4.2](#)

Changes in 5.4.1

The following features and enhancements were available in IBM Spectrum Control 5.4.1

Monitoring snapshot information for volumes that are protected by Safeguarded Copy

With Safeguarded Copy in DS8000® 8.5.0 or later, you can improve cyber resiliency by frequently creating protected, point-in-time backups of critical data, with minimum impact and effective resource utilization.

You can now view snapshot information for volumes that are backed up using the Safeguarded Copy feature across all your monitored DS8000 storage systems. Use this information to monitor the volume capacity that is protected and the capacity that is used to store volume backups for Safeguarded Copy.

Viewing topology and location information for DS8000 Fibre Channel (FC) ports

You use IBM Spectrum Control to monitor DS8000 storage systems, and you need to understand the port topology and the location of the ports in the hardware hierarchy? Not a problem. From the Resources menu, click Block Storage Systems, double-click a DS8000 storage system, and click FC Ports.

You can now check whether an FC port uses the FICON® or the SCSI FCP protocol and view the frame, I/O enclosure, and host adapter for the port.

Enhanced platform support

The following platforms are now supported:

- Db2 11.5.4
- IBM Cognos® Analytics 11.1.7
- Red Hat® Enterprise Linux 8 and Windows 2019 are now supported for all components of IBM Spectrum Control.

[View more information about the changes in 5.4.1](#)

Changes in 5.4.0

The following features and enhancements were available in IBM Spectrum Control 5.4.0:

Connecting directly to Brocade switches and fabrics

You can now connect directly to Brocade switches that run Fabric OS 8.2.1 or later. You no longer need Brocade Network Advisor (BNA) or extra servers to monitor your switches and fabrics. Connect directly to the switch by using the Fabric OS REST API and get the same key inventory and performance information that you got when you used BNA.

Monitoring performance for Hitachi Virtual Storage Platform F and G Series

You can now monitor the performance of Hitachi Virtual Storage Platform F and G Series storage systems. You can also define performance alerts so that you're notified of bandwidth, latency, and other issues before they impact your storage environment.

Viewing the name and VOLSER properties for DS8000 CKD volumes

For DS8000, you can now view both the volume name and the volume serial number (VOLSER) for your count-key-data (CKD) volumes.

User roles for monitoring IBM Spectrum Virtualize

You can now use a role with a lower level of authority on storage systems that run IBM Spectrum Virtualize 8.3.1.2 or later to collect performance metadata. When you add these storage systems, users with a lower level of authority than Administrator, such as users with the Monitor role, can collect performance metadata.

Name updates for licenses

The following licenses were renamed:

Names of IBM Spectrum Control 5.3 licenses	Names of IBM Spectrum Control 5.4.0 licenses
IBM Spectrum Control Standard Edition	IBM Spectrum Control
IBM Spectrum Control Standard Select Edition	IBM Spectrum Control Select Edition

Standard Edition bonus: If you owned a *Standard* license in a previous release, you now have access to the following features:

- Chargeback, consumer, and rollup reports
- Storage reclamation

Advanced Edition: If you own the Advanced Edition, you are not impacted by the renaming of product licenses.

Enhancements to the Licensing page

The Licensing page (Settings > Licensing) was updated to make it easier for you to understand your licensing and determine how much storage you can add before exceeding your license.

The following updates were made:

- The automatic assignment of storage category for most storage systems. Fewer manual assignments are now required.
- Improved notifications to make you aware of actions that you must complete.
- The names of column headings were changed and the hover help was revised to help you complete your compliance check.
- The capacity to be licensed is now shown in tebibytes (TiB) for consistency with the other capacity values on the page.

[View more information about the changes in 5.4.0](#)

Changes in 5.3.7

The following features and enhancements were available in IBM Spectrum Control 5.3.7:

Setting capacity limits

Set capacity limits for storage systems and pools. You want to know how much capacity you have left before your storage systems or pools are 80% full? Just set a capacity limit and you'll know how close you are to reaching your capacity limit. Then, check the Capacity-to-Limit (GiB) value and see how much capacity you can use before you reach your limit.

Support for Pure Storage devices

Monitor Pure FlashArray//M and FlashArray//X. View information about the capacity, space usage, and performance of this non-IBM storage. Other features, such as alerting, health monitoring, advanced analytics, and reporting are also supported.

More support for IBM FlashSystem® devices

You can now view information about the capacity, space usage, and performance of the FlashSystem 5000, FlashSystem 5100, FlashSystem 7200, and FlashSystem 9200 storage systems.

Monitoring DS8000 capacity on flash drives

Monitor the volume capacity that Easy Tier® places on Tier 1 and Tier 2 flash, high-capacity drives on DS8000 storage systems. View the capacity and available capacity of Tier 1 and Tier 2 flash drives in a pool and the distribution of volume extents across each of the Easy Tier drive classes.

Monitoring compliance with your license

For storage systems that use the capacity license model, view the capacity breakdown of the storage environment by the categories of drives that are used. In this release, Category 1 was added for Storage Class Memory (SCM) drives and managed disks on Spectrum Virtualize for Public Cloud.

[View more information about the changes in 5.3.7](#)

Changes in 5.3.6

The following features and enhancements were available in IBM Spectrum Control 5.3.6:

Monitoring performance for Dell EMC Unity and NetApp storage systems running ONTAP 9

Support for monitoring and alerting on performance is now added for the following non-IBM storage systems:

- Dell EMC Unity
- NetApp storage systems that are running ONTAP 9

Monitoring Hitachi Virtual Storage Platform F Series and G Series

Add Hitachi Virtual Storage Platform (VSP) F Series and G Series to the list of non-IBM storage systems that you can monitor and report on directly in IBM Spectrum Control.

More support for IBM FlashSystem devices

You can now view information about the capacity, space usage, and performance of the FlashSystem 5000, FlashSystem 5100, FlashSystem 7200, and FlashSystem 9200 storage systems.

Gaining insights into capacity usage and capacity growth

When you log in to IBM Spectrum Control, you want to know how much capacity you have, how much capacity you've used, and how much usable capacity is still available. To make this information more consumable, the overview charts have got a makeover. But, that's not all. Now, with a click, you can access key capacity metrics for each pool and volume in your storage systems. And, to help you plan capacity, two new metrics were added:

- Recent Fill Rate (%)
- Recent Growth (GiB)

Aligning capacity terms

To enhance your experience of our products and to provide a unified and simplified view of capacity concepts, IBM Storage is aligning the capacity terms across all IBM Storage products. To ensure that we use the same capacity term for the same capacity concept across our products, IBM Spectrum Control has changed 47 of the capacity terms that it uses, removed 4 terms that are no longer needed, and added 1 new term. [Learn more about the capacity terminology changes.](#)

Viewing capacity savings for devices that support hardware compression

You can now view the amount and percentage of capacity that is saved for devices that use inline data compression technology, such as IBM FlashSystem 9100 and IBM Storwize® V7000 Gen3.

Monitoring the cache fullness of pools on IBM Spectrum Virtualize

See how full your cache is by adding new cache fullness metrics to the performance charts when you view the performance of pools in your storage systems. Add and track cache fullness metrics to identify the pools that are experiencing heavy cache usage. You can also use the metrics to help investigate problems with volumes in a pool and to determine whether to move a volume to a different I/O group where the pool's cache partition does not have a cache fullness problem.

Monitoring the aggregated cache fullness of nodes at storage system level on IBM Spectrum Virtualize

See how full your nodes cache is at storage system level by adding new cache fullness metrics to the performance charts when you view the performance of storage systems. Add and track cache fullness metrics to identify the nodes that are experiencing heavy cache usage. You can also use the metrics to help investigate problems for volumes in a pool and to determine whether to move a volume to a different I/O group where the pool's cache partition does not have a cache fullness problem.

Monitoring unmap operations at host connection level on IBM Spectrum Virtualize

Track the workload of unmap operations that each host contributes to the system. To do this, you can measure the performance of unmap volume operations at host connection level. Metrics, such as the average number of MiBs that are unmapped from volumes, I/O rates, data rates, and response times for unmap operations, are collected and shown.

Monitoring Storage Class Memory (SCM) on IBM Spectrum Virtualize

You can now monitor the volume capacity that Easy Tier places on SCM drives on IBM Spectrum Virtualize systems, such as IBM FlashSystem 9100, IBM FlashSystem 7200, and the IBM Storwize family.

Creating inventory reports about IP ports

You can now create and schedule predefined reports about the IP ports on the nodes in IBM Spectrum Virtualize systems, such as SAN Volume Controller, IBM FlashSystem 9100, FlashSystem V9000, and the IBM Storwize family.

Adding IP ports to general groups

You can now add IP ports to general groups so that you can receive alert notifications about changes in the configuration or attributes of the IP ports in a general group.

Monitoring disk utilization for RAID arrays on DS8000 storage systems

For RAID arrays on DS8000 8.5.0 or later, the disk utilization values were updated in performance charts and reports. You can now more accurately check how busy the disks in the array are over a period and identify arrays that are underutilized and overutilized.

Support for IBM Cognos Analytics

IBM Spectrum Control now supports IBM Cognos Analytics Version 11.1.4. Note that IBM Cognos Analytics *does not* run on Windows Server 2019 and Red Hat Enterprise Linux 8.

[View more information about the changes in 5.3.6](#)

Changes in 5.3.5

The following features and enhancements were available in IBM Spectrum Control 5.3.5:

Important capacity updates

Capacity values were updated to better reflect capacity usage in devices that support hardware compression. Capacity, such as pool capacity and allocated space, is now measured based on the physical capacity of the device, rather than the effective capacity.

The devices affected by this update are:

- Storage systems such as FlashSystem 9100, FlashSystem 900, and Storwize V7000 Gen3, which contain IBM FlashCore® Modules with hardware compression.
- Storage virtualizers such as SAN Volume Controller and Storwize family that use back-end storage systems with hardware compression.

The effective capacity information is still available. Additional columns were added to the Block Storage Systems and Pools pages to show the effective capacity based on the data reduction savings that are achieved with hardware compression. Check the column names that begin with "Effective", for example, Effective Capacity (GiB) and Effective Used Capacity (%).

This update also affects alert thresholds and historical charting for storage systems and pools. In historical capacity charts, you'll see a sharp drop in the values for capacity and used space.

[Learn more about the capacity updates.](#)

Product fixes and security updates

In every release of IBM Spectrum Control, IBM fixes issues that are reported by our internal teams and by our customers. Along with fixes, improving the operation and security of the product is also a priority. See what's been updated and fixed in this release:

- [Fixes that are included in this release](#)
- [Security bulletins for IBM Spectrum Control](#)

[View more information about the changes in 5.3.5](#)

Changes in 5.3.4

The following features and enhancements were available in IBM Spectrum Control 5.3.4:

Monitoring unmap operations on IBM Spectrum Virtualize

You can now measure the performance of unmap volume commands and define alerts so that you are notified when unmap operations values fall outside thresholds you specify. Metrics, such as the average number of MiBs that are unmapped from volumes, I/O rates, data rates, and response times for unmap operations, are collected and shown:

- At the pool level
- At the node level
- At the I/O group level

For more information, see [Performance metrics for resources that run IBM Spectrum Virtualize](#).

Enhanced OS support

You can now install IBM Spectrum Control and Storage Resource agents on servers that run Windows Server 2019 and Red Hat Enterprise Linux 8.

Product fixes and security updates

In every release of IBM Spectrum Control, IBM fixes issues that are reported by our internal teams and by our customers. Along with fixes, improving the operation and security of the product is also a priority. See what's been updated and fixed in this release:

- [Fixes that are included in this release](#)
- [Security bulletins for IBM Spectrum Control](#)

[View more information about the changes in 5.3.4](#)

Changes in 5.3.3

The following features and enhancements were available in IBM Spectrum Control 5.3.3:

IBM Storage Insights for IBM Spectrum Control

IBM Storage Insights for IBM Spectrum Control is an IBM Cloud® service that can help you predict and prevent storage problems before they impact your business. It is complementary to IBM Spectrum Control and is available at no additional cost if you have an active license with a current subscription and support agreement for IBM Virtual Storage Center, IBM Spectrum Storage Suite, or any edition of IBM Spectrum Control.

For more information, see [IBM Storage Insights for IBM Spectrum Control](#).

Monitoring IBM Storwize V5000E storage systems

You can now monitor Storwize V5000E storage systems. You can view information about the capacity, space usage, and performance of the storage systems. Other features, such as alerting, health monitoring, advanced analytics, and reporting are also supported.

Viewing application performance

You can now view the total workload of an application that is consuming storage resources. Previously, to investigate the workload of an application, you had to review the workloads of the individual volumes. Now you can select an application from the Applications page and click View Performance to see the Total I/O Rate and the Total Data Rate for the application. On the applications details page, you can view the Aggregated Workload charts.

Faster ways to view and acknowledge your alerts

Now it's easier and faster to view and acknowledge your alerts. On the Alerts dashboard, just double-click an alert to see details of the alert on the same screen, including a performance chart for performance alerts. You can view summary details of several alerts at the same time, and acknowledge them all with a single click.

To filter alerts, you can now click one of the alert category tags. For example, to show critical and warning unacknowledged alerts, click Critical and Warning. To show your acknowledged alerts, click Acknowledged.

Automate changing the password for Spectrum Control on Windows

You can now run the `changepassword.bat` script from Windows Command Line Interface, or from an automated script, to change the passwords for the Spectrum Control user ID or for the database user ID.

Improved support for special characters in Spectrum Control and Db2 passwords

IBM Spectrum Control has improved the support for special characters in passwords, giving you the flexibility that you need based on your security requirements.

IBM Spectrum Control ports have TLS 1.1 and 1.0 disabled by default to increase security

Prior to 5.3.3, IBM Spectrum Control ports had TLS 1.1 and 1.0 enabled by default. To increase security IBM Spectrum Control ports have TLS 1.1 and 1.0 disabled by default. Therefore, IBM Spectrum Control won't be able to communicate with devices that do not support TLS 1.2. If you want to upgrade your devices to a version that supports TLS 1.2, contact your vendor.

[View more information about the changes in 5.3.3](#)

Changes in 5.3.2

The following features and enhancements were available in IBM Spectrum Control 5.3.2:

Configuring alerts using alert policies

Configuring alerts is now easier! You can now use alert policies to configure alerts for groups of resources rather than just for individual resources.

Simply create an alert policy and specify the alerts that you want, and those alerts are automatically applied to all the resources that you include in the policy.

Modifying alerts for resources is also a snap because the changes that you make are automatically applied to all the resources in the policy at the same time.

For more information about alerts and alert policies, see [How alerts work](#).

Reporting

- Create and schedule predefined reports about the capacity that is assigned to hypervisors.
- Create an inventory report about your servers' ports showing information about the ports that your servers are connected to.
- Easily generate predefined capacity reports about the resources that you are interested in. It just takes a click to select managed disks by storage systems, pools by storage systems, servers by applications, volumes by servers.
- Send custom and predefined reports as CSV and HTML attachments. Recipients can then download and open the report in a tool of their choice.
- Can't wait for the next run of a monthly report? Just click and run the report now to share the information with your colleagues!

For more information about reporting, see [Reporting](#).

Monitoring compression savings for RAID arrays in IBM FlashSystem 900, model AE3 storage system

For FlashSystem 900, model AE3 with firmware 1.6, you can now view the amount of capacity that is saved when the flash modules in the RAID array use inline data compression.

Identifying mirrored volume relationships on storage systems that run IBM Spectrum Virtualize

To help you identify the primary and secondary copies in mirrored volume relationships, the following properties for volumes were renamed:

- Target is now named Mirror Role and shows the value "Primary" or "Secondary".
- Mirror is now named Copy ID.

Tip: The name changes for these columns are also reflected in the volume information that is retrieved by using the REST API for IBM Spectrum Control. Ensure that you update any 3rd-party tools or reports that you use to show information from the REST API.

[View more information about the changes in 5.3.2](#)

Changes in 5.3.1

The following features and enhancements were available in IBM Spectrum Control 5.3.1:

- Reporting enhancements:
 - From a single page, create predefined capacity reports that you can schedule, send by email and save to your file system.
 - Create predefined inventory reports for your storage systems and for the storage resources that they use, such as nodes.
 - Add filters to custom reports. For example, you want to create a report that shows whether your file systems have enough capacity.
- You can now monitor Storwize V7000 Gen3 storage systems. You can view information about the capacity, space usage, and performance of the storage systems. You can also view the savings that are achieved when data is reduced, compressed, and deduplicated.

[View more information about the changes in 5.3.1](#)

Changes in 5.3.0

The following features and enhancements were available in IBM Spectrum Control 5.3.0:

- Create custom reports. From any table view, create reports to share information with your colleagues and managers about the capacity, configuration, and health status and the performance of the resources that you monitor in your storage environment.
- Monitor IBM FlashSystem 9100 storage systems.
- Prevent your FlashSystem 900 from running out of capacity by monitoring the physical capacity and receiving alerts when the physical capacity is low. View information about pool capacity and the used and available capacity in the RAID array.
- Add multiple storage systems for monitoring at the same time.

[View more information about the changes in 5.3.0](#)

Changes in 5.2.17

The following features and enhancements were available in IBM Spectrum Control 5.2.17:

- View detailed information about the enclosures for block storage systems that run IBM Spectrum Virtualize, such as FlashSystem V9000, FlashSystem 900, SAN Volume Controller, and the IBM Storwize family.
- Use new overview charts for storage systems to answer key physical capacity, volume capacity, and capacity savings questions:
 - How much space is available.
 - How much of your thin-provisioned volume space is used.
 - How much space you are saving by using compression, deduplication, and thin provisioning.
- For storage systems that run IBM Spectrum Virtualize 8.1.1 or later, use the Pool Attributes column to distinguish data reduction pools from standard pools and identify if the data in the pool is encrypted. For storage systems that run IBM Spectrum Virtualize 8.1.3 or later, you can also view the capacity savings that are made when you use data reduction pools.
- For storage systems that run IBM Spectrum Virtualize 8.1.3 or later, you can see the amount and percentage of capacity that is saved when data is deduplicated. The capacity savings for deduplication are shown at the storage systems, pools, and volumes levels.
- For storage systems that run IBM Spectrum Virtualize, you can track the historical capacity and the capacity savings that were achieved when volumes are compressed. This allows you to monitor the efficiency of compression over time.
- Use the Compressed and Deduplicated filters to show only volumes with the selected filter value.

[View more information about the changes in 5.2.17](#)

Changes in 5.2.16

The following features and enhancements were available in IBM Spectrum Control 5.2.16:

- Monitor the availability of clusters by checking which nodes have spare nodes. You can also check the status of the spare nodes, such as which nodes are online and which nodes are on standby.
- Monitor the workload of compressed volumes on storage systems that run IBM Spectrum Virtualize™. You can view I/O rate, data rate, and response time metrics for compressed volumes and uncompressed volumes. See [Compression metrics for volumes](#).
- In previous releases, IBM Cognos Business Intelligence reports was installed with and could be accessed directly from IBM Spectrum Control. In this release, IBM Cognos Analytics is installed as a separate product. See [View more information about the changes in 5.2.16](#).
- The integration with external LDAP repositories in IBM Spectrum Control is improved.

[View more information about the changes in 5.2.16](#)

Changes in 5.2.15

The following features and enhancements were available in IBM Spectrum Control 5.2.15:

- Creating agentless servers used to be a manual process, and automating it was a frequently requested enhancement. IBM Spectrum Control now automatically generates agentless servers for the physical servers and virtual machines in your SAN environment, so that you can easily see how they are consuming storage.
 - Change the default names of the tiers that are shown in the GUI.
 - Add or change properties, such as the location or custom tags, for multiple servers, storage systems, hypervisors, switches, or fabrics. You can use the properties to filter or sort the resources in the GUI or in an external application if the data is shared or exported.
 - Define more than 30 new alerts for volumes in block storage systems, including alerts for capacity values such as Compression Savings, Data Deduplication Savings, Data Reduction Savings, and Unallocated Volume Space. Use these alerts to detect and monitor potential configuration and capacity issues in your environment.
 - Find volumes with the same attributes. Use the new filters, such as Easy Tier and Mirror, to show only volumes with the selected filtered value. You can also sort the volumes by allocated and unallocated space, and much more.
 - Add and monitor IBM Spectrum Virtualize software-only clusters.
 - IBM Db2 Version 11.1 is now included with IBM Spectrum Control.
 - Use the Representational State Transfer (REST) API to automatically export the data that you need to generate custom reports or share with external applications.
- The following information is now available:
- How much capacity is used by applications, departments, hypervisors, and physical servers.
 - How much block capacity is used by a specific application, department, hypervisor, or physical server.

[View more information about the changes in 5.2.15](#)

Changes in 5.2.14

The following features and enhancements were available in IBM Spectrum Control 5.2.14:

- Enhancements for IBM Cloud Object Storage:
 - Add COS Vaults to applications and general groups so you can track how much space is being consumed by logically related vaults and how that consumption trends over time.
 - Use the enhanced Vaults by Accessers chart, and the new Vaults by Drives and SliceStors by Drives charts, to monitor the availability of all your COS Vaults and COS SliceStor® nodes.
 - View the number of failed drives in the COS SliceStor nodes and the tolerance of the nodes to more drive failures.
 - View the tolerance of COS Vaults to drive failures across all the COS SliceStor nodes in the storage pool.
 - View the COS Vaults, access pools, and storage pools that are used by IBM Spectrum Scale storage systems.

- The user experience is enhanced within the GUI for integration with an external LDAP repository. IBM Spectrum Control is now available to a larger set of users and groups and you are able to log in to IBM Spectrum Control with one set of credentials.
- IBM Spectrum Control ensures, when it sends alert notifications and reports by email to your email server, that it complies with the security standards that you configure on your email server.
- You can now choose the unit of measurement for a capacity metric when you create or edit chargeback and consumer reports. So, you can generate reports that show storage consumption in the unit of measurement of your choice whether it is PB or PiB, TB or TiB, or GB or GiB.
- View the following port delay metrics for storage systems that run IBM Spectrum Virtualize: Port Delay Time, Port Delay I/O Percentage
- DS8000 storage systems, you can now see the hosts and host clusters that host connections belong to.
- View a list of all the primary volumes that are not protected by a copy data relationship.
- View which Managed Disks are active Quorum disks.
- View the site name of each node in a stretched cluster.
- For Storwize storage systems, 7.8 or later, you can now view information about flash drives by tier in storage pools and volumes. For storage pools, you can view the capacity and available capacity for solid-state drives (Tier 0) and for read-intensive flash storage (Tier 1). For volumes, you can view the capacity for solid-state drives (Tier 0) and for read-intensive flash storage (Tier 1).
- Upload logs automatically to IBM Software Support to help resolve problems quickly.

[View more information about the changes in 5.2.14](#)

Changes in 5.2.13

The following features and enhancements were available in IBM Spectrum Control 5.2.13:

- View more information about the capacity, space usage, and performance of Dell EMC. Support for versions of Dell EMC SMI-S Provider or Dell EMC Solutions Enabler that are compliant with SMI-S 1.6 provides a more representative view of Dell EMC VMAX, Dell EMC VNX, and Dell EMC VNXe storage systems.
- Gain insights into the external storage that is used by IBM Spectrum Scale:
 - View capacity values for the external storage that clusters and file systems use. This view provides insight into the total amount of data that is migrated from all file systems in a storage system, or the amount of data that is migrated from specific file systems.
 - Check whether you have enough internal space to recall data that was migrated to external storage.
- For storage systems that run IBM Spectrum Virtualize, you can now see the cluster that host connections belong to.
- Enhancements for IBM Cloud Object Storage:
 - View information about the hard disk drives, solid-state drives, and flash modules in COS Slicestor nodes.
 - In the Vaults by Accessers chart, view the vaults that cannot be accessed and the vaults that are at risk of access failure.
- On the new VDisk Mirrors and HyperSwap® pages in the Copy Data panel, you can analyze storage device relationships across your block storage environment in IBM Spectrum Control to identify how redundancy is affecting available capacity.
- You exclude the cost of storage in chargeback and consumer reports when you want to highlight storage consumption and not storage cost.

[View more information about the changes in 5.2.13](#)

Changes in 5.2.12

The following features and enhancements were available in IBM Spectrum Control 5.2.12:

- Monitor the capacity, space usage, and performance of FlashSystem A9000 and FlashSystem A9000R storage systems.
- Create consumer reports that show the capacity and the cost of the block storage that an application, department, hypervisor, or physical server uses.
- Report on allocated versus assigned capacity in chargeback and consumer reports.
- Analyze replication and FlashCopy® relationships across your block storage environment
- View the amount of external storage from cloud services and other storage providers that is being used by the file systems in a IBM Spectrum Scale cluster. For external storage that is provided by IBM Cloud Object Storage, you can also view capacity information, including the percentage of space that is being used.
- Monitor sites, mirrors, and vault quotas for IBM Cloud Object Storage. For example, you can:
 - Monitor the capacity and space usage of sites. See which COS Slicestor and COS Accesser® nodes are located at each site and which sites cannot tolerate COS Accesser node failures.
 - View the space quotas that are configured for vaults and see whether any vaults are violating the soft or hard quota limits.
 - View information about mirrors, such as the pair of vaults in each of the mirrors, and the number of COS Accesser nodes that are configured to access the mirrors. You can also see the storage pools and sites for the vaults in the mirror.
- Define alerts that notify you of status changes on the IBM Spectrum Scale nodes that are configured as cloud gateways.
- Customize the analysis period for analyzing tiering and balancing pools to the days of the week and hours of the day of your choice.
- Access information about your resources by using the Representational State Transfer (REST) API for IBM Spectrum Control. You can use this information to generate custom configuration and performance reports.

[View more information about the changes in 5.2.12](#)

Changes in 5.2.11

The following features and enhancements were available in IBM Spectrum Control 5.2.11:

- Changes to alerts:
 - Create alerts for a set of storage resources that are assigned to an application.
 - Define alerts for a set of resources in custom groups that you create.
 - Combine multiple alert conditions from multiple resources in a single, custom alert.
- Create, schedule, and send chargeback reports by email to make the owners and managers aware of the cost of the capacity and the amount of capacity that is used by their applications, departments, hypervisors, and servers.
- View key performance metrics for IBM® FlashSystem 900 storage systems:
 - Response times
 - I/O rates and data rates
 - The condition of the storage system
- Review new capacity charts to detect capacity shortages and investigate trends in storage growth for your tiers.
- Analyze replication relationships across your block storage environment.
- Monitor the status, capacity, and space usage of IBM Cloud Object Storage.
- Identify which file systems in a Spectrum Scale cluster use external storage that is provided by cloud services or other storage providers.

- Enable automated probe scheduling to simplify the process of scheduling probes and help avoid excessive load on the product server.
- Determine whether the total capacity of the storage systems that IBM Spectrum Control manages exceeds your product license entitlement.
- When you complete a fresh installation of the product, the number of random listening ports is reduced. Information about the ports is documented and you can change the Device server ports after you complete the installation or perform an upgrade.
- Manually uninstall IBM Tivoli® Storage Productivity Center for Replication from IBM Spectrum Control servers on your computer.
- IBM Db2® Version 10.5 Fix Pack 7 is now included with IBM Spectrum Control.

[View more information about the changes in 5.2.11](#)

Changes in 5.2.10

The following features and enhancements were available in IBM Spectrum Control 5.2.10:

- Identify key performance metrics that are outside of a standard range on IBM System Storage® SAN Volume Controller and IBM Storwize® storage systems.
- Export performance data for a resource to a compressed file.
- Review new capacity charts to detect capacity shortages and investigate trends in storage growth for your block, file, and object storage resources.
- Review the capacity and space usage that the tiered storage in your environment consumes. You can see when your tiered storage will run out of space and the weekly growth rate in storage usage for each tier.
- View the values for the capacity of virtualized storage and non-virtualized storage to understand how capacity is used when storage virtualization is implemented in your storage environment.
- Use rollup reporting to combine capacity and status information from multiple instances of IBM Spectrum Control™. You can use this information to gain a network-wide perspective of storage usage in your environment when you have multiple IBM Spectrum Control servers deployed.
- Group resources, such as the storage systems with lease agreements that end in the current year, so that you can view information about the resources at one location.
- IBM Spectrum Control V5.2.10 now supports only Storage Management Initiative (SMI) data sources, also called CIM agents, for managing switches from Brocade Communications Systems, Inc.
- View the resource types that were explicitly added to an application or subcomponent.

[View more information about the changes in 5.2.10](#)

Changes in 5.2.9

The following features and enhancements were available in IBM Spectrum Control 5.2.9:

- Reclaim unused storage capacity and more effectively allocate volume space on your block storage systems.
- View the amount of storage that can be reclaimed in each of the storage systems that are monitored.
- Use the volume reclamation recommendations to reclaim unused storage space and use your storage more efficiently. By reclaiming the volumes, you can reduce costs by recycling existing storage space instead of purchasing new storage media.
- IBM Spectrum Control™ now displays the switch blade number, port number, and blade names just as they were defined in other tools, such as Brocade Network Advisor.
- IBM Db2® Version 10.5 Fix Pack 6 is now included with IBM Spectrum Control.

[View more information about the changes in 5.2.9](#)

Changes in 5.2.8

The following features and enhancements were available in IBM Spectrum Control 5.2.8:

- The name of the product was changed from IBM Tivoli Storage Productivity Center to IBM Spectrum Control.
- Comprehensive improvements to alerting capabilities:
 - You can now define alerts for all the key attributes of a resource, including attributes for status, configuration, capacity, and performance.
 - To avoid receiving too many emails and disruptive notifications for some alerts, you can now choose a new suppression option where only one notification is sent.
 - If you want to define more than one alert for the same attribute but with different criteria and settings, you can duplicate that alert.
- View the performance of clusters, nodes, and file systems in IBM Spectrum Scale.
- Monitor object storage systems on IBM Spectrum Scale.
- Monitor the capacity and space usage of IBM FlashSystem 900 storage systems.
- View charts that show the current and historical capacity of your block storage pools.
- IBM Spectrum Control is now a native, 64-bit application for all server operating systems.
- Simple Network Management Protocol (SNMP) Version 3 is now the preferred protocol for the management of switches and fabrics from Cisco Systems, Inc., and is an optional protocol for devices from Brocade Communications Systems, Inc. and other vendors.
- View information about the performance of inter-switch connections.
- View information about Brocade switches that are in Access Gateway mode. You can also view information about the routes from ports that are connected to a hypervisor, server, or storage system to ports that are connected to fabrics.
- View information about Brocade F_port trunks and about NPV links.
- You can now more easily identify the internal resources that are encountering problems in your environment.
- Compare how much capacity is managed by IBM Spectrum Control with the total amount of capacity that you are entitled to manage.
- View information about the capacity for each of the class drives (SSD, Nearline HDD, and Enterprise HDD), that Easy Tier can use to tier or distribute volume extents.
- To consolidate the functions of IBM Spectrum Control into a single GUI, the stand-alone GUI was removed.

[View more information about the changes in 5.2.8](#)

Capacity terminology changes

IBM Storage is making the capacity terminology that is used in IBM Storage products consistent. In Q1 2020, the capacity terminology was made clearer and simpler across IBM Storage products, including IBM Spectrum Control.

Discontinued features in IBM Spectrum Control

View a list of the features that were discontinued in different releases of IBM Spectrum® Control.

What is a discontinued feature?: Discontinued features are features that are no longer supported in IBM Spectrum Control. In some cases, these features were removed from the GUI and are no longer available. In other cases, these features might still appear in the GUI, but they are no longer supported and might not work as intended. For both cases, it's recommended that you use another feature or product offering to complete the same actions, when possible.

Discontinued feature	Discontinued in version	Resolution
Db2® 11.1 is no longer supported as the database repository for storing metadata about your monitored devices.	5.4.7	Upgrade Db2 to Db2 11.5 or later. For information out how to upgrade Db2, see Example of upgrading Db2 in a Windows environment .
Operating systems where you can install Storage Resource agents: <ul style="list-style-type: none"> HP-UX 11i v3 Oracle Solaris 10 and 11 	5.4.0	Install your Storage Resource agents on different servers with supported operating systems. For a list of supported operating systems, see https://www.ibm.com/support/pages/node/6249361#Agents .
Storage systems: <ul style="list-style-type: none"> DS3000, DS4000®, DS5000, or DS6000™ Non-IBM storage systems that are monitored with SMI-S 1.1 providers. NetApp devices that run ONTAP 8. Huawei devices 	5.4.0	No future support is planned for monitoring these storage systems. For a complete list of storage systems that are supported in 5.4.0 and later, see IBM Spectrum Control interoperability matrix for storage systems .
Provisioning of block and file storage systems, which includes: <ul style="list-style-type: none"> Cloud configuration concepts, such as service classes and capacity pools Reserved capacity concept External Application role 	5.4.0	No other feature in IBM Spectrum Control currently replaces this feature. For your future provisioning and automation needs with IBM storage, consider using another product, such as Ansible® by Red Hat®. For more information about Ansible, see https://www.ansible.com/integrations/infrastructure/ibm-storage .
Optimization: <ul style="list-style-type: none"> Tiering Balancing Volume transformation 	5.4.0	No other feature in IBM Spectrum Control currently replaces this feature. It's recommended that you use a different product to optimize tiering, balance storage, and transform volumes. Pro tip: With the increasing use of all flash storage and environments with a few large pools rather than many smaller pools, optimization and balancing functions are becoming less critical.
Snapshot-based protection that is provided by IBM Spectrum Protect™ Snapshot for the applications and databases that IBM Spectrum Control uses.	5.4.0	This feature is still available in the IBM Storage Virtual Center storage solution.
STAR data is no longer available in the following: <ul style="list-style-type: none"> Cognos® Analytics reporting tool CLI (tpctool) commands IBM Spectrum Control GUI 	5.4.0	<ul style="list-style-type: none"> For your reclamation needs, go to Advanced Analytics > Reclamation. For tier-related capacity management, go to Groups > Tiers. For other uses of your STAR data, it's recommended that you use a different product.
Tivoli® Storage Productivity Center for Replication	5.3.0	Use IBM Copy Services Manager. See IBM Copy Services Manager .
IBM Spectrum Control workflows for storage provisioning with IBM Tivoli Provisioning Manager	5.3.0	No other feature in IBM Spectrum Control currently replaces this feature.
The agent for integrating IBM Spectrum Control alerting with IBM Tivoli Monitoring	5.3.0	No other feature in IBM Spectrum Control currently replaces this feature.
OpenStack Cinder driver that uses the provisioning feature in IBM Spectrum Control	5.3.0	Use the Cinder driver that is associated with your storage system.
Database views for custom reporting with SQL	5.3.0	Use the Representational State Transfer (REST) API to access information about resources and to generate custom configuration and performance reports. See IBM Spectrum Control REST API .
The tpctool CLI commands for provisioning that don't use the service-class concepts of IBM Spectrum Control	5.3.0	No other feature in IBM Spectrum Control currently replaces this feature.
Using IBM Spectrum Control components to enable IBM DS8000® storage systems to work with external authentication servers, such as LDAP	5.3.0	No other feature in IBM Spectrum Control currently replaces this feature.
vSphere Web Client extension for IBM Spectrum Control and VASA 1.0 provider	5.3.0	No other feature in IBM Spectrum Control currently replaces this feature. However, IBM Spectrum Connect provides IBM Storage Enhancements for VMware vSphere Web Client and an IBM Storage Provider for VMware VASA. For more information, see the IBM Spectrum Connect Knowledge documentation .
Graphical view of the data path between storage resources, SAN components, and servers (Data Path Explorer)	5.3.0	To view the relationships between resources: <ol style="list-style-type: none"> Go to a resource list page, such as the Block Storage Systems page. Right-click a resource and select View Details. On the details page, view the Related Resources section. Click a related resource to view more information about it. Related resources that you can view include servers, file systems, fabrics, switches, object storage systems, and back-end storage systems for storage virtualizers.

Sponsor user program

Sponsor users interact directly with designers and developers to improve the user experience and to help shape the future of the overall storage portfolio.

IBM® needs your experience and expertise as an active participant to work with our designers and developers to create new features for the future and help improve the overall look and feel of our IBM Storage products.

Are you ready to align us with your reality of what a product should do?

Are you a real user or a potential user of an IBM Storage software product?

IBM wants you to be excited about being an active participant and establishing a collaborative and fluid relationship between your company and our design and development teams.

Express® your thoughts and get your ideas heard! Then, as the discussion moves forward, we will integrate your insights and feedback into our decision making process and make them action items. In addition, as an active program user you get a peek into upcoming storage releases and gain beta access to new products.

Note: If you are interested in joining the IBM Storage Sponsor User Program or just want to get additional information, complete the following form: [IBM Storage Sponsor User](#).

Important: Always remember your information is kept confidential and is only used by the IBM Design and Development for product development purposes.

Beta program

The IBM Spectrum® Control Beta is a continuous program. It gives you a first look at upcoming features, a chance to influence design, an opportunity to test the new features in your own environment, and a direct voice into the product development process.

There are many reasons why new and existing customers, IBM® Business Partners, and even current members of IBM become part of the Beta program as a tester.

Customer quote: Bob Oesterlin, Sr. Principal Storage Engineer at Nuance: "If you're a IBM Spectrum Control user, definitely participate in IBM's Beta program. It gives you access to the latest features early on, direct access to developers and a great community of other early adopters. I've been a part of the IBM Spectrum Control Beta program for years, and the IBM team has been fantastic and responsive."

Benefits

Some benefits of IBM's Beta program include:

Gain access to early code and evaluate new product features and enhancements

You get access to the Beta code before general availability of the product release to determine whether the new features and enhancements are a good fit for your business organization. Once the code is downloaded, you are able to validate and test the new software directly in your environment. You can then identify and then fix any concerns before the code is available saving you precious time and preventing any production issues later on. When the code is made available, you are ready for installation, implementation and to take advantage of the new capabilities.

Interact with design and development

The product designers, architects, developers, and testers are integrated into the Beta program to help support the participants. They can assist you with any issues that you might encounter. You can also request an advocate to work with you during the process and post questions on the Beta forum.

Collaborate with other Beta customers

The Beta program includes group meetings that provide you with an opportunity to interact with other program participants about your configuration and testing experiences with the beta features. Participants are encouraged to share their experiences with the development team.

Participate in product education

The Beta program provides the participants with education on the new features and functions that are available in a Beta. The education is usually done in presentation form on a web conference where you can get a head start in learning about the capabilities.

Become an IBM reference customer

After your positive Beta experience, IBM invites you to participate in the reference program. The IBM Marketing team helps you craft a message to let other potential Beta testers know about your success when you adopt and use beta features.

Get enrolled!

To enroll in the Beta program, complete the [IBM Spectrum Control Beta Sign up form](#).

Collaborating with the team

Collaborate with the IBM Spectrum® Control team to help improve the product.

About this task

Got a great idea for making IBM Spectrum Control even better? Do you want to vote for an enhancement that was requested by another user? The IBM® System Storage® Ideas Portal is a place where you can collaborate with the development team for IBM Spectrum Control and other users through your ability to search, view, comment on, submit and track product requests. Help shape the future of IBM!

Procedure

1. Go to the IBM System Storage Ideas Portal at <https://ibm-sys-storage.ideas.ibm.com/ideas?project=SPCO>.
2. To submit an enhancement request, click Add A New Idea.
3. Log in with your IBMid.

4. In the Choose a workspace for this idea drop-down menu, select IBM Spectrum Control .
5. Complete the form and click Add Idea.

Printable documentation


The documentation for IBM Spectrum® Control is available in PDF format.

Viewing PDF files

Click one of the following links to view the corresponding PDF within your browser.

- [IBM Spectrum Control Quick Installation Guide](#) (about 1200 KB)
- [IBM Spectrum Control Installation Guide](#) (about 3100 KB)
- [IBM Spectrum Control Administrator's Guide](#) (about 2200 KB)
- [IBM Spectrum Control User's Guide](#) (about 10200 KB)

Downloading Adobe Acrobat Reader

To view or print these PDF files, you need Adobe Acrobat Reader. You can download a free copy from the [Adobe website](https://get.adobe.com/reader/)  (<https://get.adobe.com/reader/>).

Product licenses

Learn about the product licenses that are available for IBM Spectrum® Control. Current licenses are IBM Spectrum Control and IBM Spectrum Control Select Edition. The IBM Spectrum Control Advanced Edition license is still applied, if you owned it in a previous release.

IBM Spectrum Control offers two licenses based on the type of metric that you use:

- The IBM Spectrum Control license is based on the total number of storage capacity units (SCUs) that are required so that all of your capacity is licensed for all of your storage systems. This license applies if you currently have the IBM Spectrum Control or IBM® Virtual Storage Center license.
- The IBM Spectrum Control Select Edition license (and the IBM Virtual Storage Center for Storwize® license) is based on the total number of licenses that you need for your storage systems that use enclosures.

Note: Only one difference exists between the IBM Spectrum Control and IBM Spectrum Control Select Edition licenses: the license metric (SCUs versus enclosures). The features are otherwise identical in both solutions. To view the features, see [Product overview](#).

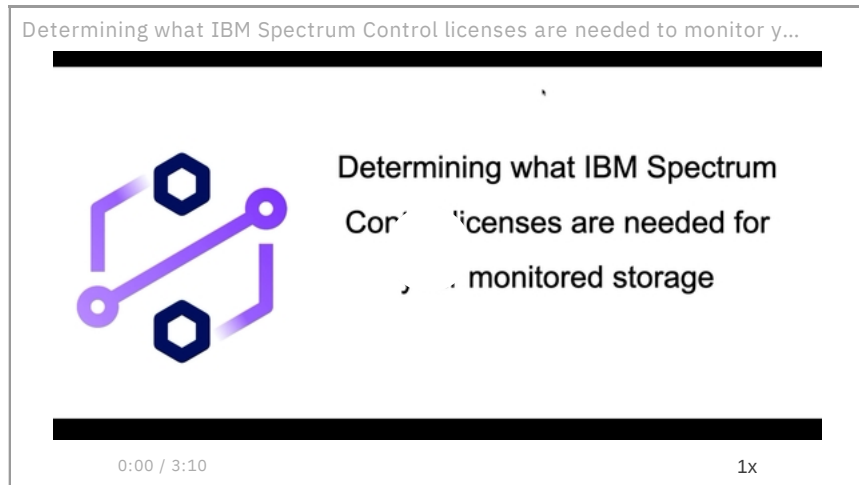
To view a complete list of devices that can be used with IBM Spectrum Control, go to <https://www.ibm.com/support/pages/node/388393>.

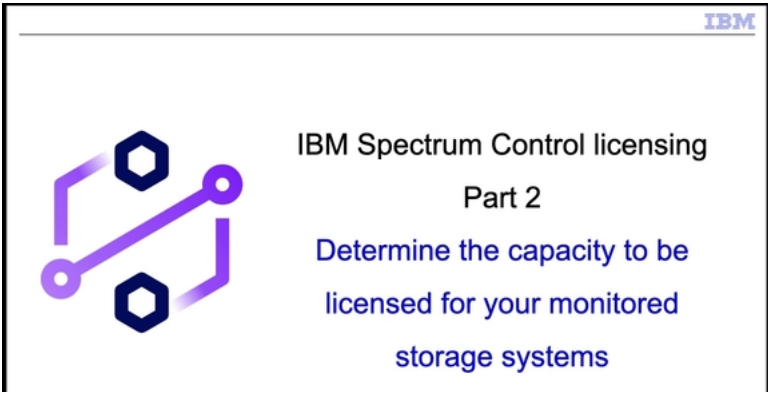
IBM Storage Insights for IBM Spectrum Control: IBM Storage Insights for IBM Spectrum Control is an IBM Cloud® service that can help you predict and prevent storage problems before they impact your business. It is complementary to IBM Spectrum Control and is available at no additional cost if you have an active license with a current subscription and support agreement for IBM Virtual Storage Center, IBM Spectrum Storage Suite, or any edition of IBM Spectrum Control.

For more information and to view a list of features in the service, see [IBM Storage Insights for IBM Spectrum Control](#).

Determining which licenses are needed for your monitored storage

Watch two short videos about how to view the capacity of the storage systems that IBM Spectrum Control monitors in your environment. Then, check the records of the licenses that you purchased to determine how much additional storage can be monitored before you're no longer covered.





IBM Spectrum Control licensing
Part 2
Determine the capacity to be licensed for your monitored storage systems

0:00 / 3:52 1x

- [Definitions of licensing metrics](#)
Learn more about the licensing metrics that are available for IBM Spectrum Control.
- [Names of equivalent licenses in previous releases](#)
IBM Spectrum Control licenses were known by different names in previous releases. Use this table to determine the names.
- [Actions that are available based on role](#)
Your IBM Spectrum Control role and product license determine the actions that are available in the product.
- [Licensing](#)
View the capacity of the storage systems that IBM Spectrum Control monitors in your environment. Then, check the records of the licenses that you purchased to determine how much additional storage can be monitored before you're no longer covered.

Definitions of licensing metrics

Learn more about the licensing metrics that are available for IBM Spectrum® Control.

About this task

Three different license metrics are available for IBM Spectrum Control:

SCUs

Licenses are priced by the number of storage capacity units (SCUs) that you need for the capacity to be licensed for the storage system. The number of SCUs is based on the capacity and categories of drives that are used by the storage system. This metric is available for the IBM Spectrum Control and IBM® Virtual Storage Center perpetual licenses.

Total TiB

Licenses are priced by the total usable capacity in tebibytes (TiB) for all storage systems that use the capacity license model. This metric is available for the IBM Virtual Storage Center subscription license.

Total TiB (with limits)

Licenses are priced by the total usable capacity in tebibytes (TiB) for all storage systems that use the capacity license model, with the following limits:

- Up to 1,500 TiBs of usable capacity
- No limit I/O groups

This metric is available for the IBM Virtual Storage Center Entry and IBM Spectrum Storage Suite licenses.

Enclosures

Licenses are priced by the number of drive slots and capacity in the storage enclosures. This metric is available for the IBM Spectrum Control Select Edition and IBM Virtual Storage Center for Storwize® licenses.

For each enclosure, the number of licenses is calculated by using the following rules. Calculations are rounded up to the nearest whole number.

1. The number of slots in the enclosure divided by 25.
2. The sum of the physical capacity of the disks in the enclosure divided by 500 TiB.
3. The higher number from rule 1 and rule 2 is the number of licenses that are needed for the enclosure.

For examples of using the enclosure pricing model, see [Calculating the number of enclosure licenses for your storage systems](#).

To learn about the licensing of IBM Spectrum Control, click the Announcement in [Table 1](#). Information about licenses is available in the section *Charge Metrics*.

Table 1. Product announcements for IBM Spectrum Control licenses

License	Announcement	Revised Announcement
IBM Spectrum Control - Storage Capacity Unit	5.3.0	5.3.6.0
IBM Spectrum Control - Select/Enclosure license	5.3.0	5.3.7.0
IBM Virtual Storage Center (subscription license)	5.4.3	

Names of equivalent licenses in previous releases

IBM Spectrum® Control licenses were known by different names in previous releases. Use this table to determine the names.

Table 1. IBM Spectrum Control licenses and their previous names

Names of IBM Spectrum Control 5.4.0 licenses	Names of IBM Spectrum Control 5.3 licenses
IBM Spectrum Control	IBM Spectrum Control Standard Edition
IBM Spectrum Control Select Edition	IBM Spectrum Control Standard Select Edition

Actions that are available based on role

Your IBM Spectrum® Control role and product license determine the actions that are available in the product.

Users who are assigned the Administrator role or the Monitor role can use product functions. The actions that are available for each function depend on the role that is assigned to the user:

Administrator role

Users who are assigned the Administrator role have access to all monitoring and administrative actions.

Monitor role

Users who are assigned the Monitor role can view information about monitored resources and other objects such as tasks, alerts, and service classes. They can acknowledge alerts and resource statuses, open logs, and open management GUIs.

The following table outlines the actions that are available only for the Administrators role. All other actions are available to the Monitor and Administrator roles. In addition to the restrictions listed in this table, users who are assigned the Monitor role do not have access to user management functions.

Table 1. Product actions that are available only to users with the Administrator role

Function	Actions that require the Administrator role
Single dashboard view of the storage environment that you can use to manage storage systems, hypervisors, servers, and Fibre Channel fabrics.	<ul style="list-style-type: none"> • Adding and removing resources • Administering connections • Scheduling data collection • Changing and viewing the automated probe schedule • Viewing and editing history retention settings • Modifying license settings
Performance monitoring for storage systems and Fibre Channel networks.	<ul style="list-style-type: none"> • Scheduling performance monitors • Starting or stopping performance monitors
Capacity and usage monitoring of resources.	<ul style="list-style-type: none"> • Scheduling probes • Starting or stopping probes • Modifying Storage Resource agents • Enabling automatic zoning
Health and alerting for hypervisors, networks, servers, and storage systems.	<ul style="list-style-type: none"> • Creating, modifying, and deleting alert policies • Setting which alert policy manages a resource • Adding and modifying resources for management by an alert policy • Defining and modifying alert definitions • Editing alert notification settings
Capacity and performance of the storage that applications, departments, and general groups use.	<ul style="list-style-type: none"> • Creating applications, departments, and general groups • Creating, modifying, and removing filters to add resources to applications • Adding and removing resources in applications and general groups, directly • Adding applications as members of other applications • Adding departments to other departments • Adding applications to departments
Storage reclamation	Viewing volumes that can be reclaimed

Function	Actions that require the Administrator role
Roll-up reporting, in which capacity data is combined from multiple instances of IBM Spectrum Control for reporting purposes.	<ul style="list-style-type: none"> Adding and removing subordinate servers Starting a probe for a subordinate server Modifying the connection information for a subordinate server
Predefined Reports <ul style="list-style-type: none"> Predefined capacity reports allow users to quickly create reports about capacity anomalies and shortfalls, which can be scheduled and sent by email or saved to the user's file system, or both. Predefined inventory reports allow users to quickly create reports about their storage resources, which can be scheduled and sent by email or saved to the user's file system, or both. 	<ul style="list-style-type: none"> Creating, deleting, and editing reports Configuring the email server Emailing reports Saving reports to the file system
Custom reports From any table view in the GUI, custom reports can be created, which can be scheduled and sent by email or saved to the user's file system, or both, about capacity of storage resources, the configuration and attributes of storage resources, and the performance of storage resources.	<ul style="list-style-type: none"> Creating, deleting, and editing reports Configuring the email server Emailing reports Saving reports to the file system
Chargeback and consumer reports <ul style="list-style-type: none"> Chargeback reports show the capacity and the cost of the storage that is used by applications, departments, hypervisors, and physical servers. Consumer reports show the capacity and the cost of the block storage that is used by an application, department, hypervisor, and physical server. 	<ul style="list-style-type: none"> Creating, deleting, and editing reports Configuring the email server Emailing reports
Capacity limits for block storage systems and pools If your company has a policy to set a limit on the capacity that is used, you can set a capacity limit. When the capacity limit is set, you can then monitor the amount of capacity that is available before the capacity limit is reached.	<ul style="list-style-type: none"> Setting capacity limits Defining alerts for capacity limits Removing capacity limits

Licensing

View the capacity of the storage systems that IBM Spectrum® Control monitors in your environment. Then, check the records of the licenses that you purchased to determine how much additional storage can be monitored before you're no longer covered.

To determine how many licenses are needed for the capacity and enclosures that you monitor, complete these steps on the Settings > Licensing page:

1. [Know your licenses.](#)
2. [Understand the licensing models.](#)
3. [Select the license model for each storage system.](#)
4. [Update any settings that were not automatically updated or correctly identified for each storage system.](#)
5. [Determine if you're covered.](#)

1. Know your licenses

To get a list of the licenses that you purchased, complete one or more of these actions:

- Contact the person within your organization who originally purchased the licenses.
- Contact your IBM® seller or IBM Business Partner.
- Go to [Passport Advantage® Online](#).

2. Understand the licensing models

The licenses that you purchased determine which license models and values on the Licensing page are used to determine if you're covered. In some cases, a combination of values is used. For example, if you have some storage systems that are licensed by capacity and some that are licensed by enclosure, both the Total SCU and the Total Enclosure License values might be used.

License	License model
IBM Spectrum Control	Measured by the number of storage capacity units (SCUs) that are needed to cover the total usable capacity of the storage systems that IBM Spectrum Control monitors. For more information, see the description of the SCU value on the Licensing page .
IBM Virtual Storage Center (perpetual license)	
IBM Virtual Storage Center (subscription license)	Measured by the total usable capacity in terabytes (TiB) of your storage. For more information about the subscription option, see IBM Virtual Storage Center .
IBM Spectrum Control Select Edition	Measured by the number of licenses that are needed for the storage enclosures, modules, or expansions that IBM Spectrum Control monitors. See Calculating the number of enclosure licenses for your storage systems .
IBM Virtual Storage Center for Storwize®	

License	License model
IBM Virtual Storage Center Entry	Measured by the total capacity in terabytes (TiB) that IBM Spectrum Control monitors.
IBM Spectrum Storage Suite	

3. Select the license model for each storage system

You must specify the license model for each of your storage systems. The license model, capacity or enclosure, is the type of IBM license that you purchased for the storage system.

License model	Description
Capacity	The category of storage that is used by the storage system is needed to calculate the SCU usage. The category is based on the drive classes of the RAID array and is automatically assigned following a probe. For IBM Spectrum Scale and virtual nodes in IBM Cloud Object Storage, you must specify the storage category.
Enclosure	A probe automatically detects the number of enclosures that the storage system manages and determines the number of enclosure licenses that are required. You can manually update the number of enclosure licenses. After you modify the number of enclosure licenses, that number will not change unless you manually update it again. Probes that detect a different number of enclosure licenses do not modify a number that was manually updated.

Tip: For IBM System Storage® DS8000®, the license model is automatically assigned after a probe.


4. Update settings for each storage system

Update any settings that were not automatically updated or correctly identified for your storage systems. Examples of these settings include the number of enclosures in a storage system that were not identified or the drives in a storage system that were not categorized.

Learn more about configuring license settings: For more information about how to update the license settings for your storage systems, see [Configuring license values for storage systems](#).

5. Determine if you're covered

To verify that your purchased licenses cover the capacity and enclosures that are monitored by IBM Spectrum Control, compare the values in your purchased licenses to the values at the top of the page.



- Look at Total TiB for these licenses: IBM Virtual Storage Center Entry Edition and IBM Spectrum Storage Suite.
- Look at Total SCU for these licenses: IBM Spectrum Control, IBM Virtual Storage Center, and older IBM Spectrum Control licenses that are based on SCU.
- Look at Total Enclosure Licenses for these licenses: IBM Spectrum Control Select Edition, IBM Virtual Storage Center for Storwize, and older IBM Spectrum Control licenses that are based on enclosure.

Learn more about license values: For more information about how capacity and enclosure values are calculated for storage systems and storage virtualization systems (such as IBM SAN Volume Controller), see [How licensing values are calculated](#).

What's next

If the Total TiB, Total SCU, and Total Enclosure Licenses values are less than the value of your purchased licenses, you're covered.

If these values are more than the value of your licenses, you can reduce the amount of storage that you're monitoring or contact your IBM seller or IBM Business Partner to purchase additional storage.

Information about the summary values

The following, summary information is shown at the top of the Licensing page:

Capacity License Model

The number in parentheses shows the total number of storage systems that are licensed by storage capacity.

Total TiB

The sum of the capacity to be licensed for all storage systems that use the capacity license model. This value applies if you have the IBM Virtual Storage Center Entry or IBM Spectrum Storage Suite license.

Total SCU

The total number of storage capacity units that you need for the capacity to be licensed for all storage systems that use the capacity license model. This value applies if you have the IBM Spectrum Control or IBM Virtual Storage Center license.

Category 1, 2, 3, or 4

The capacity breakdown of the storage environment by the categories of drives that are used by the storage systems. The breakdown is shown both as a bar chart and as terabytes (TiB).

Category Info

For storage systems that are licensed by capacity rather than enclosure, a link that describes the categories of drives that are used by RAID arrays.

Enclosure License Model

The number in parentheses shows the total number of storage systems that are licensed by the enclosures that IBM Spectrum Control monitors.

Total Enclosure Licenses

The total number of licenses that are required for storage systems that use the enclosure license model. This value applies if you have the IBM Spectrum Control Select Edition or IBM Virtual Storage Center for Storwize license.

Information about storage system values

The following information is shown for each storage system on the Licensing page:

License Model

The type of IBM Spectrum Control license that was purchased for the storage system. Systems that use the capacity model are charged by the capacity to be licensed or the SCUs. Systems that use the enclosure model are charged by the number of enclosure licenses.

SCU

The number of storage capacity units that the license must cover for the storage system.

An SCU is the measure of capacity by which the license for a storage system is charged. The number of SCUs is based on the drive classes that are used by the storage system and the capacity that is assigned to each category. The following table shows the SCU categories, drive classes, and ratios that are supported by IBM Spectrum Control.

Table 1. Supported SCU licenses, drive classes, and ratios

Category	Drive Classes	Ratio of usable storage to SCU
1	Storage Class Memory (SCM) drives and managed disks on IBM Spectrum Virtualize for Public Cloud	1 TiB to 1 SCU
2	Flash and Solid-State Drives (SSDs)	1 TiB to 0.847 SCU 1.18 TiB to 1 SCU
3	10K Serial Attached SCSI (SAS) Drives, 15K Fibre Channel Drives, and storage systems that use Category 4 drives with advanced architectures	1 TiB to 0.5 SCU 2 TiB to 1 SCU
4	Near Line SAS (NL-SAS), Serial ATA (SATA) Drives, and Safeguarded Copy (SGC) Target Capacity	1 TiB to 0.25 SCU 4 TiB to 1 SCU
Tips: <ul style="list-style-type: none">Any storage capacity that uses drive classes that are not listed in the table is classified as Category 1.Calculations are rounded up to the nearest whole number.		

The following table provides examples that you can use when you want to determine your SCU usage.

Table 2. Examples of SCU usage

Storage category	TiB usable	SCU ratio	Total SCUs
1	42	1	42
2	400*	1.18	339
3	800	2	400
4	1600	4	400
Totals	2842		1181

Category

For capacity license models, the breakdown of the storage environment by the four categories of drives that are used by the storage system. The breakdown is shown as tebibytes (TiB).

For most storage systems, the storage category is set automatically. However, you must set the category for IBM Spectrum Scale and virtual appliances in IBM Cloud Object Storage. If the storage category is not set, IBM Spectrum Control assumes, for the calculations of SCU usage, that Category 1 drives are used.

For enclosure license models, the number of licenses that you need for the storage enclosures that the storage system manages. You can modify the number of enclosure licenses. For example, a probe might detect that two enclosure licenses are needed for a storage system, but you're aware of two external enclosures that were not detected. You calculate that these enclosures need one license each. In that case, you can change the probe-detected value (2) to the more accurate number of enclosure licenses (4).

Keep in mind the following restrictions before you modify the number of enclosure licenses:

- You cannot reduce the number of enclosure licenses to be less than the number that was detected by a probe.
- After you modify the number of enclosure licenses, that number will not change unless you manually update it again. Probes that detect a different number of enclosure licenses do not modify a number that was manually updated.

Restriction:

- For SAN Volume Controller, if a probe doesn't detect any enclosures, 1 is displayed automatically.
- For XIV® systems and IBM Spectrum Accelerate, the number of enclosures can't be edited.

Capacity to be Licensed (TiB)

The capacity that must be licensed for the storage system. Capacity that is reserved for formatting or RAID management is not included. For storage virtualizers and IBM Spectrum Scale, this value includes the capacity that is provided by back-end storage systems that are not being monitored by IBM Spectrum Control.

Capacity from back-end storage systems that are monitored by IBM Spectrum Control is recorded separately on the Licensing page and is not included in this value.

Monitored Back-end Capacity (TiB)

For storage virtualizers and IBM Spectrum Scale, the capacity that is provided by back-end storage systems that are being monitored by IBM Spectrum Control.

This capacity is not included in the capacity to be licensed of the virtualizer or IBM Spectrum Scale. Each monitored back-end storage system has a separate record on the Licensing page.

Unmonitored Back-end Capacity (TiB)

For storage virtualizers, the capacity that is provided by back-end storage systems that are not being monitored by IBM Spectrum Control. This capacity is included in the capacity to be licensed of the virtualizer. Unmonitored storage systems do not have their own record on the Licensing page.

Total Capacity (TiB)

The total capacity of the storage system, including all monitored and unmonitored back-end storage.

$$\text{Total Capacity} = \text{Capacity to be Licensed} + \text{Monitored Back-end Capacity}$$

Uncategorized Capacity (TiB)

The total capacity that is not categorized for the storage system. Capacity that is not categorized is also included in the Capacity to be Licensed value.

For most storage systems, IBM Spectrum Control automatically assigns the storage category for the pools on the storage system, based on the technology type of the drives in the pool. However, for IBM Spectrum Scale and virtual appliances on IBM Cloud Object Storage, IBM Spectrum Control cannot determine the storage category.

To ensure that the correct storage category is used, assign the storage category on the Licensing page. If the storage category is not set, IBM Spectrum Control assumes, for the calculations of SCU usage, that Category 1 drives are used.

Enclosure Licenses

The number of enclosure licenses that you need for the storage system.

- [How licensing values are calculated](#)
Learn about how the capacity to be licensed and the number of enclosure licenses that are needed for your storage systems are calculated.
- [Configuring license values for storage systems](#)
Discover the capacity that IBM Spectrum Control monitors so that you can determine whether you are in compliance with your license.
- [Licensing examples](#)
See how the capacity to be licensed and other values on the Licensing page are calculated for different storage systems and configurations.

How licensing values are calculated

Learn about how the capacity to be licensed and the number of enclosure licenses that are needed for your storage systems are calculated.

Calculating the number of enclosure licenses for your storage systems

For each enclosure, the number of licenses that are needed is calculated by using the following rules. Calculations are rounded up to the nearest whole number.

1. The number of slots in the enclosure divided by 25.
2. The sum of the physical capacity of the disks in the enclosure divided by 500 TiB.
3. The higher number from rule 1 and rule 2 is the number of licenses that are needed for the enclosure.

Examples of the enclosure licensing calculations.

Rule 1 - Number of slots in the storage enclosures (or modules or expansions) that are managed:

- Enclosure with 12 slots requires 1 license.
- Enclosure with 92 slots requires 4 licenses.

Rule 2 - Sum of the physical capacity of the disks in the enclosure:

- Enclosure with 200 TiB requires 1 license.
- Enclosure with 1,250 TiB requires 3 licenses.

Example of a storage system with 4 enclosures:

Enclosure	Slots	Rule 1 - Licenses required	Capacity (TiB)	Rule 2 - Licenses required	Licenses required
Enclosure 1	28	2	400*	1	2
Enclosure 2	92	4	3,700	8	8
Enclosure 3	60	3	1,750	4	4
Enclosure 4	16	1	225	1	1
Total enclosure licenses required					15

Calculating the capacity to be licensed for your storage systems

Capacity to be licensed is the usable capacity of the storage systems that are being monitored and that is made available for storage consumption. Apart from IBM Spectrum Scale and IBM® Cloud Object Storage, the capacity to be licensed is based on the usable capacity of the RAID arrays on the storage system.

Information about the calculation of capacity to be licensed:

- Capacity that is reserved for overheads, such as RAID management is not included.
- Data reductions savings at the storage system level, such as compression and deduplication, are not included. For storage systems that use data reduction technologies, the capacity to be licensed is the usable capacity, not the effective capacity.
- Capacity to be licensed is the total capacity of the storage system after RAID is applied and spare disks and capacity are deducted. So, capacity to be licensed is the capacity that can be provisioned to servers, when no over-provisioning is used. Any form of cache, for example, DRAM, SSD, or flash storage is not included, and the capacity to be licensed is less than the raw capacity of the storage system.

Refer to the following information for details about how capacity to be licensed is calculated for different storage systems:

Storage virtualization systems

Capacity to be licensed is calculated for the following storage virtualizers:

- Storage systems that run IBM Spectrum Virtualize, such as IBM SAN Volume Controller and IBM FlashSystem® 9100.
- Non-IBM storage systems such as Hitachi VSP.

A storage virtualizer includes capacity from other storage systems. Capacity from back-end storage systems that are not monitored by IBM Spectrum® Control is counted in the capacity to be licensed of the virtualizer. Capacity from back-end storage systems that are monitored by IBM Spectrum Control is not counted. The

license for the monitored back-end storage is recorded separately and has its own record on the Licensing page.

For storage virtualizers that have internal capacity, such as IBM Storwize® V7000, FlashSystem 7200, and IBM FlashSystem 9100, the capacity to be licensed is calculated as:

**sum of the physical capacity of the internal RAID arrays +
sum of the capacity of the external MDisks from unmonitored back-end systems**

For storage virtualizers that do not have internal capacity, such as some models of IBM SAN Volume Controller, the capacity to be licensed is calculated as:

sum of the capacity of the external MDisks from unmonitored back-end systems

Block and file storage systems

The capacity to be licensed is the sum of the physical capacity of the RAID arrays on the storage system.

Exception: For FlashSystem A9000, FlashSystem A9000R, and IBM XIV® Storage System, the capacity to be licensed is the overall physical capacity of the storage system.

For the following storage systems, the capacity and enclosure license information is not calculated because IBM Spectrum Control does not collect the RAID array information from the device. The value **Unavailable** is shown in the SCU and Capacity to be Licensed columns on the Licensing page.

- Block and file storage systems that are managed by SMI-S providers. For example:
 - NetApp Data ONTAP 8.1
 - Dell EMC VMAX, VNX, and VNXe
 - Hitachi VSP F series

To get a complete picture of your licensing requirements, complete the following tasks for these storage systems:

- If the storage system is licensed by enclosure, you must manually enter the number of enclosure licenses. The value is then automatically included in the Total Enclosure Licenses value in the summary area at the top of the page.
- If the storage system is licensed by capacity, you must manually calculate the SCUs that you need for the storage system. Then, add that value to the Total SCU value in the summary area.

IBM Spectrum Scale storage systems

The capacity to be licensed is the sum of the capacity of the Network Shared Disks (NSDs) that are assigned to the file system pools.

NSD capacity from back-end storage systems that are not monitored by IBM Spectrum Control is included.

The following capacity is not included in the capacity to be licensed:

- NSDs that are not assigned to a pool.
- NSD capacity from back-end storage systems that are monitored by IBM Spectrum Control. The license for the monitored back-end storage is recorded separately and has its own record on the Licensing page.
- NSD capacity on external pools.

For IBM Spectrum Scale storage systems that are configured for file and object storage, the capacity to be licensed for the object storage is included in the file system pool capacity and is not counted separately.

Don't forget to assign the storage category for your file system pools: The storage category is based on the technology type of the drives in the pool, for example, Storage Class Memory and Flash, and is used in the calculation of SCU usage. To ensure that the correct storage category is used, assign the storage category on the Licensing page.

If the storage category is not set, IBM Spectrum Control assumes, for the calculations of SCU usage, that Category 1 drives are used.

Calculating enclosure licenses for IBM Spectrum Scale: For IBM Spectrum Scale storage systems that are licensed by enclosure, you need one enclosure license for each NSD server node that is connected to NSDs.

IBM Cloud Object Storage systems

The capacity to be licensed is the sum of the capacity of the COS Slicestor® nodes in IBM Cloud Object Storage.

Don't forget to assign the storage category for your virtual COS Slicestor nodes: The storage category is based on the technology type of the drives in the node, for example, Storage Class Memory and Flash, and is used in the calculation of SCU usage. A COS Slicestor node can be a server or a virtual appliance. For the server nodes, IBM Spectrum Control automatically assigns the Near-Line drive category. For virtual appliances, IBM Spectrum Control cannot determine the storage category of the drives in the node.

For the calculation of SCU usage for the node, IBM Spectrum Control assumes that Category 1 drives are used in the virtual appliance. To ensure that the correct storage category is used in the calculations, assign the storage category on the Licensing page for each virtual node.

Calculating enclosure licenses for IBM Cloud Object Storage

To calculate the number of enclosure licenses, the same rules that are specified in [Calculating the number of enclosure licenses for your storage systems](#) are used.

For the calculation:

- Each COS Slicestor node is treated as an enclosure.
- Instead of enclosure slots, the calculation uses the number of drives in the node.
- The capacity is based on the total capacity of the drives in the node.

Example of IBM Cloud Object Storage with 3 COS Slicestor nodes:

Nodes	Drives	Rule 1 - Licenses required	Capacity (TiB)	Rule 2 - Licenses required	Licenses required
Node 1	16	1	400	1	1
Node 2	64	3	2,400	5	5
Node 3	32	2	350	1	2
Total enclosure licenses required					8

Configuring license values for storage systems

Discover the capacity that IBM Spectrum® Control monitors so that you can determine whether you are in compliance with your license.

About this task

To determine the storage capacity that IBM Spectrum Control monitors, you must confirm the type of IBM license that you purchased for each storage system (known as the *license model*).

If you do not know the type of IBM Spectrum Control license that you purchased for each storage system, you can skip the process and return to complete the configuration later.

Whether you log on to IBM Spectrum Control for the first time or after an upgrade, the procedure for configuring license compliance is similar.

Procedure

1. Click Settings > Licensing.
 - If you are logging on to IBM Spectrum Control for the first time, the Set License Model page is displayed, with no selections made in the License Model column. Go to step 2.
 - If you are logging on to IBM Spectrum Control after an upgrade, the Licensing page is displayed, with the selections from the previous release in the License Model column. If some of your storage systems were not classified by license model and pool category in the last release, you must classify them by clicking Assign licenses.
Important: For some storage systems that support the IBM® Easy Tier® IBM feature, the capacity is not calculated when you upgrade from IBM Spectrum Control. In that case, it is necessary to probe the storage systems again. Following an upgrade and before you open the Licensing page for the first time, be sure to run a probe of all of your monitored storage systems.
2. Choose the type of license model. For storage systems that are licensed by enclosure, enter the number of enclosure licenses that you need for the storage system.
Tip: A storage enclosure that is managed by IBM Spectrum Control is an independently powered, channel-attached device that stores data on magnetic disks or solid-state drives. An example is a disk controller and its expansion units. Each expansion unit is a separate enclosure. It can be the main controller that houses disk or solid-state drives, or the expansion chassis that houses extra disk or solid-state drives to expand the total capacity of the storage system.
3. The IBM Spectrum Control software automatically assigns the category of disk for the pools:
 - If all automatic assignments are correct, click Finish.
 - If one or more assignments are incorrect, click the correct category for each pool and click Finish.

When you finish entering information, the Licensing page is displayed. The page presents data about the storage in your environment. If the page shows uncategorized storage systems, you can click the Assign licenses link to enter its license model and drive category. The page also indicates the total tebibytes (TiB) of capacity that IBM Spectrum Control manages. Values are shown as tebibytes (TiB). 1 TiB is equal to 2^{40} (1,099,511,627,776) bytes. To determine license compliance, compare the total TiB with your proof of entitlement that you received when you purchased your licenses.

Licensing examples

See how the capacity to be licensed and other values on the Licensing page are calculated for different storage systems and configurations.

In these examples, SVC-1 is an IBM® SAN Volume Controller and is monitored by IBM Spectrum® Control.

Example 1: Back-end capacity is not monitored

This example uses the following configuration:

- SVC-1 has no internal storage and does not receive capacity from storage systems other than the ones used in this example.
- Flash-1 is a Storwize® V5000, which manages one enclosure with a drive capacity of 750 TiB. The drive capacity is configured as a 600 TiB RAID array.

Here's the situation:

- Flash-1 is not monitored by IBM Spectrum Control.
- Flash-1 provides 300 TiB of capacity to SVC-1.

The following values are shown on the Licensing page:

Storage System	Enclosure Licenses	Capacity to be Licensed (TiB)	Monitored Back-end Capacity (TiB)	Unmonitored Back-end Capacity (TiB)	Total Capacity (TiB)
SVC-1	0	300	0	300	300

Explanation:

The 300 TiB Flash-1 capacity that is provided to SVC-1 is shown as unmonitored back-end capacity and is included in the capacity to be licensed for SVC-1.

The remaining 300 TiB capacity on Flash-1 is not provided to storage systems that are monitored by IBM Spectrum Control and does not need to be licensed.

Example 2: Back-end capacity is monitored

Configuration:

- SVC-1 has no internal storage and does not receive capacity from storage systems other than the ones used in this example.
- Flash-1 is a Storwize V5000, which manages one enclosure with a drive capacity of 750 TiB. The drive capacity is configured as a 600 TiB RAID array.

Situation:

- Flash-1 is monitored by IBM Spectrum Control and is licensed by capacity.
- Flash-1 provides 300 TiB of capacity to SVC-1.
- SVC-1 receives an additional 800 TiB of capacity from other storage systems that are not monitored by IBM Spectrum Control.

The following values are shown on the Licensing page:

Storage System	Enclosure Licenses	Capacity to be Licensed (TiB)	Monitored Back-end Capacity (TiB)	Unmonitored Back-end Capacity (TiB)	Total Capacity (TiB)
Flash-1	0	600	0	0	600

Storage System	Enclosure Licenses	Capacity to be Licensed (TiB)	Monitored Back-end Capacity (TiB)	Unmonitored Back-end Capacity (TiB)	Total Capacity (TiB)
SVC-1	0	800	300	800	1100

Example 3: Back-end capacity is monitored and the storage system is licensed by enclosure

Configuration:

- SVC-1 has no internal storage and does not receive capacity from storage systems other than the ones used in this example.
- Flash-1 is a Storwize V5000, which manages one enclosure with a drive capacity of 750 TiB. The drive capacity is configured as a 600 TiB RAID array.

Situation:

- Flash-1 is monitored by IBM Spectrum Control and is license by enclosure.
- Flash-1 provides 300 TiB of capacity to SVC-1.

The following values are shown on the Licensing page:

Storage System	Enclosure Licenses	Capacity to be Licensed (TiB)	Monitored Back-end Capacity (TiB)	Unmonitored Back-end Capacity (TiB)	Total Capacity (TiB)
Flash-1	2	600	0	0	600
SVC-1	0	0	300	0	300

Explanation:

- There is no capacity to be licensed for SVC-1. The back-end capacity that is provided to SVC-1 is licensed on Flash-1.
- To calculate the number of enclosure licenses, the Flash-1 drive capacity (750 TiB) is used, rather than the RAID array capacity (600 TiB).

Calculation of Enclosure Licenses for Flash-1

The higher number from rule 1 and rule 2 is used.

- Rule 1: 24 slots in the enclosure, which requires 1 enclosure license
- Rule 2: 750 TiB drive capacity, which requires 2 enclosure licenses (750 / 500, rounding up to the nearest whole number)





For more information, see [Calculating the number of enclosure licenses for your storage systems](#).

- Because Flash-1 is licensed by enclosure, the capacity to be licensed is not included in the Total SCU and Total TiB values at the top of the Licensing page.

Getting started with IBM Spectrum Control

Learn about the key tasks for setting up, installing, and monitoring IBM Spectrum Control to manage the resources and infrastructure in your storage environment.

Tip: Click items in the image to view information about installing, configuring, upgrading, and troubleshooting IBM Spectrum Control.

 Preparing	Understand licensing	Review release information	Explore social media
	Plan for installation	Review installable components	Install Db2
	Install IBM Spectrum Control in a single-server environment	Install IBM Spectrum Control in a multiple-server environment	Install IBM Cognos Analytics
 Installing	Open IBM Spectrum Control GUI	Add resources	Configure alert notifications
	Define alerts and alert policies	Authorize users	Configure history and data retention
	Prepare for an upgrade	Upgrade in a single-server environment	Upgrade in a multiple-server environment
 Upgrading	Migrate reports to IBM Cognos Analytics		
	Start and stop IBM Spectrum Control servers	Troubleshoot installing, uninstalling, and upgrading IBM Spectrum Control	Troubleshoot IBM Spectrum Control GUI
	Resolve problems		
 Troubleshooting			

Social media for IBM Spectrum Control

Watch videos, read blogs, and explore social media to learn more about how to use IBM Spectrum® Control to manage your storage environment. Be social! Join the conversation at [#IBMStorage](#) and [#softwaredefinedstorage](#).

- **Videos for IBM Spectrum Control**
Watch videos about how to use IBM Spectrum Control to manage your storage environment.
- **Blogs for IBM Spectrum Control**
Read blogs and follow us on Twitter to learn from the experts and get the latest tips about how to use IBM Spectrum Control to manage your storage environment.

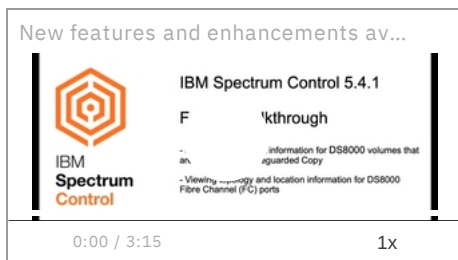
Videos for IBM Spectrum Control

Watch videos about how to use IBM Spectrum® Control to manage your storage environment.


Important: The following videos were recorded by using multiple versions of IBM Spectrum Control, but are at least partially applicable to the current version of the product.

New features and enhancements available in IBM Spectrum Control 5.4.1


Learn about the new features and enhancements to IBM Spectrum Control 5.4.1.




Part 1: Determining which IBM Spectrum Control licenses are needed for your monitored storage
Learn how to use the Licensing page to find out which licenses are needed for the capacity and enclosures that you monitor.




Part 2: Determining which IBM Spectrum Control licenses are needed for your monitored storage
Check out part 2 of the licensing videos to continue learning about how to use the Licensing page to find out which licenses are needed for the capacity and enclosures that you monitor.




Troubleshooting storage infrastructure performance
Learn how to monitor, compare, and troubleshoot performance on storage systems.




Viewing and extracting performance data into a CSV file
Learn how to extract performance data from IBM Spectrum Control using the GUI.



Monitoring the performance of IBM Spectrum Scale
Learn how to monitor IBM Spectrum Scale clusters.



Editing resource locations in bulk
Learn how to edit the locations of resources in bulk and use Cognos® BI to report on capacity by location.



Blogs for IBM Spectrum Control

Read blogs and follow us on Twitter to learn from the experts and get the latest tips about how to use IBM Spectrum® Control to manage your storage environment.

IBM Spectrum Control has a strong presence on social media. Explore blogs and forums to read the perspectives of storage management experts and gain insights about how to use IBM Spectrum Control.

- [IBM® Storage Community](#)
- Twitter:
 - [#IBMStorage](#)
 - [#ibmsystems](#)
 - [#softwaredefined](#)
 - [Documentation hints and tips](#)

Table 1. Recent blogs for IBM Spectrum Control

Blog	Summary
A new subscription option for IBM Virtual Storage Center is now available!	Learn about the new subscription option for new customers who purchase an IBM Virtual Storage Center license.
IBM Spectrum Control: Performance monitoring hints and tips	Get some hints and tips to help you use IBM Spectrum Control Target file not found to monitor the performance of your storage.
IBM Spectrum Control: Performance monitoring for FlashSystem 900 -- don't forget the SNMP agent	Learn how to fix a performance monitoring problem for IBM FlashSystem® 900.
Sign up for the free 90-day trial of IBM Spectrum Control today!	Learn how to sign up for the free 90-day trial of IBM Spectrum Control , which provides an opportunity to use a set of tools for managing storage capacity, availability, alerts, events, performance, and storage systems.
IBM Spectrum Control: Spark joy with this Pro Tip for sorting switches	Learn how to sort switches so that physical switches and their virtual or logical switches are grouped together.
New release! IBM Spectrum Control, 5.4.2 has an update for IBM Virtual Storage Center	Learn about the features that were introduced in IBM Spectrum Control 5.4.2.
Maximize your IBM DS8000® monitoring with the new release of IBM Spectrum Control!	Learn about the features that were introduced in IBM Spectrum Control 5.4.1.
Connections really matter! Direct connection to Brocade or Broadcom switches highlights new release of IBM Spectrum Control!	Learn about the features that were introduced in IBM Spectrum Control 5.4.0.
Whether Pure Flash Array monitoring or more capacity enhancements IBM Spectrum Control adds to your overall business uptime!	Learn about the features that were introduced in IBM Spectrum Control 5.3.7.
More, more and even more storage system monitoring and enhanced capacity reporting in IBM Spectrum Control 5.3.6!	Learn about the features that were introduced in IBM Spectrum Control 5.3.6.
More third-party storage support & IBM Spectrum Virtualize family enhancements highlight IBM Spectrum Control 5.3.5 release	Learn about the features that were introduced in IBM Spectrum Control 5.3.5.
IBM Storage Insights for IBM Spectrum Control IBM Cloud® service highlights IBM Spectrum Control, 5.3.3!	Learn about the features that were introduced in IBM Spectrum Control 5.3.3.
Configuring alerts with alert policies highlights new release of IBM Spectrum Control!	Learn about the features that were introduced in IBM Spectrum Control 5.3.2.
New release of IBM Spectrum Control is now available: More reporting enhancements!	Learn about the features that were introduced in IBM Spectrum Control 5.3.1.
Easy report creation and IBM FlashSystem 9100 monitoring highlight new release of IBM Spectrum Control	Learn about the features that were introduced in IBM Spectrum Control 5.3.0.
New version of IBM Spectrum Control adds more monitoring information for IBM Spectrum Virtualize storage systems	Learn how IBM Spectrum Control provides more monitoring information for IBM Spectrum Virtualize, information on migrating reports to IBM Cognos® Analytics, and integration with external LDAP repositories, among other features and enhancements.
IBM Spectrum Control virtual mailbox: Refresh my memory, where does IBM Spectrum Control keep log files for its main components?	Learn about the locations of the log files for IBM Spectrum Control. You can view these log files to identify where problems occur, and you can provide them to IBM Support so that they can help resolve any problems.
IBM Spectrum Control virtual mailbox: Is there a way to see capacity summaries for my storage systems?	Learn how to use the overview charts for storage systems to answer key physical capacity, volume capacity, and capacity savings questions.

Table 2. Blogs about IBM Spectrum Control 5.2.15

Blog	Summary
------	---------

Blog	Summary
Edit resource locations in bulk and use Cognos BI to report on capacity by location	Learn how to edit the locations of multiple resources at the same time and use Cognos BI to report on capacity by location.
A great time-to-value feature: IBM Spectrum Control auto-creates your agentless servers	Learn how IBM Spectrum Control now automatically generates agentless servers for the physical servers and virtual machines in your SAN environment, so that you can easily see how they are consuming storage.
Agentless servers : how to handle an inconsistency in your storage environment	Learn how to merge agentless servers that represent the same physical server so that a single agentless server is shown in IBM Spectrum Control.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge documentation to discover information and tips that you might not be aware of but could find useful.

Table 3. Blogs about IBM Spectrum Control 5.2.14

Blog	Summary
IBM Spectrum Control 5.2.14	Learn about the features that were introduced in IBM Spectrum Control 5.2.14.
Monitoring a stretched cluster or Hyperswap environment	Learn how to monitor your stretched cluster or Hyperswap environment with IBM Spectrum Control.
How does today's response time compare with last week's?	Learn how to compare performance metrics across different time ranges to help troubleshoot potential bottlenecks in an application.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge documentation to discover information and tips that you might not be aware of but could find useful.
IBM Spectrum Control virtual mailbox	Learn how to upload logs automatically and send them to IBM Support for investigation.

Table 4. Blogs about IBM Spectrum Control 5.2.13

Blog	Summary
IBM Spectrum Control 5.2.13	Learn about the features that were introduced in IBM Spectrum Control 5.2.13.
Chargeback and Consumer Reports: Don't show me the money!	Learn how to create a custom report that sends you a daily summary of the capacity used by applications without including storage costs.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge documentation to discover information and tips that you might not be aware of but could find useful.
IBM Spectrum Control virtual mailbox	Learn how much capacity is consumed by replication. As a storage administrator you are asked to provide a capacity report across your entire storage environment to your CTO. A critical part of capacity reporting is understanding how much capacity is consumed by replication.

Table 5. Blogs about IBM Spectrum Control 5.2.12

Blog	Summary
IBM Spectrum Control 5.2.12	Learn about the features that were introduced in IBM Spectrum Control 5.2.12.
Consumer Reports: For your eyes only!	Learn how to configure a report that shows the block capacity and cost of the block capacity for a single storage consumer. Each consumer gets to star in their own report.
Chargeback and Consumer Reports: Same price for the same storage	In chargeback and consumer reports, you want to ensure that you charge the same price for the same type of storage whether it is block storage, copy data, or file storage. And, if your block storage is tiered, you want to charge the same price for each tier to all of your storage consumers.
Chargeback and Consumer Reports: How much storage was really eaten?	When you create chargeback and consumer reports, you can charge for the block storage that is allocated to the storage consumer or the block storage that is assigned to the storage consumer.
Look at that IBM FlashSystem data reduction!	FlashSystem A9000 and FlashSystem A9000R can reduce the size of your data significantly by using technologies such as compression and data deduplication. Learn how to see exactly what sort of savings you are getting, and how that compares to savings on other storage systems.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge documentation to discover information and tips that you might not be aware of but could find useful.

Table 6. Blogs about IBM Spectrum Control 5.2.11

Blog	Summary
IBM Spectrum Control 5.2.11	Learn about the features that were introduced in IBM Spectrum Control 5.2.11.
Chargeback Reports: Who's eating my storage?	Whether you manage or own applications, departments, hypervisors, or physical servers, you want to know how much storage is being consumed and how much it costs to maintain the storage that is consumed. And, you want that information in your inbox now, every week or month so that you can keep a close eye on storage costs and capacity consumption.
Capacity and compliance: Know your software license in IBM Spectrum Control	Your business organization just purchased a software license(s) for IBM Spectrum Control or any IBM product that includes IBM Spectrum Control such as IBM® Virtual Storage Center or IBM Spectrum Storage Suite and you want to assess your storage license situation.
Show your boss how your IBM FlashSystem compares	Learn how to compare the performance of your IBM FlashSystem to other storage systems. You can compare the response times and the I/O rates of IBM FlashSystem storage systems to other storage systems.
See how Cognos BI can help you create capacity and performance reports	Learn how to use the Cognos® Business Intelligence (BI) reporting tool to report on the condition, capacity, and performance of your storage resources.
From the IBM Spectrum Control hints and tips virtual mailbox: Is there a way to simplify the process of scheduling probes?	The automated probe option enables the creation of a probe and collects status and asset information more efficiently about your storage resources for monitoring.
Pearls in the documentation	Delve into the IBM Spectrum Control Knowledge documentation to discover information and tips that you might not be aware of but could find useful.
Navigating the new knowledge documentation	Learn how to quickly find information in the new IBM Documentation and read about tips for getting the most out of its new interface.

Table 7. Blogs about IBM Spectrum Control 5.2.10

Blog	Summary
IBM Spectrum Control 5.2.10	Learn about the features that were introduced in IBM Spectrum Control 5.2.10.

Blog	Summary
Extracting bulk performance data for analysis was never easier!	Learn about how to export bulk performance data for a storage system or fabric to a compressed file. If you contact IBM Support to help you analyze a performance problem with a resource, you might be asked to send this file.
IBM Spectrum Control: Rollup reporting provides a dynamic view into a large enterprise environment	Learn about how to use rollup reporting to gain a network-wide perspective of storage usage in your environment when you have multiple IBM Spectrum Control servers deployed.

Table 8. Blogs about IBM Spectrum Control 5.2.9

Blog	Summary
IBM Spectrum Control 5.2.9	Learn about the features that were introduced in IBM Spectrum Control 5.2.9.

Table 9. Blogs about IBM Spectrum Control 5.2.8

Blog	Summary
IBM Spectrum Control 5.2.8	Learn about the features that were introduced in IBM Spectrum Control 5.2.8.
Monitoring the entire data path of object storage systems	Learn about how to monitor object storage implementation on IBM Spectrum Scale and troubleshoot along the entire data path.

Product overview

IBM Spectrum® Control provides a set of tools for managing storage capacity, availability, alerts, events, performance, and resources. It can reduce the complexity of managing a storage environment by centralizing, simplifying, and optimizing storage tasks that are associated with storage systems, storage networks, performance troubleshooting, and capacity management.

IBM Spectrum Control can help you detect potential problems on your storage devices. For example, it can notify you when a server or storage system is running out of disk space or warn you of impending storage hardware failure so that you can prevent unnecessary system and application downtime.

IBM Spectrum Control provides the following benefits and functions:

- Simplifies the management of storage infrastructures
- Manages, administers, and provisions SAN-attached storage
- Monitors and tracks performance of SAN-attached resources
- Monitors, manages, and controls (through zones) SAN fabric components
- Manages the capacity utilization and availability of file systems

Get a quick overview of the storage monitoring features that are available in IBM Spectrum Control:

Table 1. Overview of features in IBM Spectrum Control

Features	Description
Understanding your environment	A single dashboard view of your storage environment to help you manage storage devices.
	Monitor IBM and non-IBM block storage, file storage, object storage, hypervisors, fabrics, switches.
	Capacity monitoring of storage systems, pools, volumes, file systems, and file sets.
	Performance monitoring of storage systems, Fibre Channel networks, applications, and departments.
	Drill down performance workflows to troubleshoot bottlenecks.
	Define groups of resources that represent applications or departments.
	Explore storage and server virtualization relationships.
Health and alerting	Explore replication relationships.
	Show health status of hypervisors, networks, servers, and storage systems, and selected internal entities.
	Alert on status changes and thresholds of capacity, performance, and error metrics.
Reporting	Customizable, multi-conditional alerting and the ability to create alert policies for many devices.
	Inventory, capacity, and performance reports.
	Reports that you can create from the information in the user-interface tables. You can configure, schedule, and save this information to your file systems. Or you can specify to send it to another user or person by mail.
	Chargeback reports for showing the capacity and the cost of the storage that is used by applications, departments, hypervisors, and physical servers.
	Consumer reports for showing the capacity and the cost of the block storage that is used by an application, department, hypervisor, and physical server.
	Roll up reports in which capacity data is combined from multiple instances of IBM Spectrum Control for reporting purposes.
	Reports that use REST API and Cognos®.
Analytics	For storage systems that run IBM Spectrum Virtualize, the transform storage feature moves volumes, and configures compression and the properties of the volumes.
	Storage reclamation that you can realize when the block storage systems in your data center are analyzed, and you determine and reclaim volumes that aren't used for storing data. You can replace the existing storage space instead of purchasing new storage media.
	Business impact analysis (applications, departments, and groups).
General	Active directory and LDAP integration for managing users.
	Customization options, including the ability to define the history retention of performance and capacity metadata, define how you are notified of alerts and notifications, and schedule when capacity and configuration metadata is automatically collected.
	Inclusion of the license for IBM® Copy Services Manager to manage 2-site replication, 3-site replication, and advanced copy services. Learn more about how to download and install Copy Services Manager at https://www.ibm.com/docs/en/csm .
	Snapshot-based protection that is provided by IBM Spectrum Protect™ Snapshot for the applications and databases that IBM Spectrum Control uses.

- [Supported devices in IBM Spectrum Control](#)
Find out which storage devices and products that you can monitor with IBM Spectrum Control.
- [Architecture](#)
The IBM Spectrum Control consists of several components that form the infrastructure of its storage-management functions.
- [Interfaces for IBM Spectrum Control](#)
IBM Spectrum Control provides multiple user interfaces for managing the storage infrastructure in an enterprise environment.

- **[Starting IBM Spectrum Control](#)**
You can start IBM Spectrum Control by opening a web browser and entering a web address for the IBM Spectrum Control logon page. For example, you might enter `https://storage.example.com:9569/szm`.
- **[Navigation](#)**
IBM Spectrum Control provides many of the functions for managing a storage environment. To access the functions in its GUI, use the menu bar at the top of the main window.
- **[IBM Spectrum Storage Suite](#)**
IBM Spectrum Control provides monitoring, automation and analytics for multiple-vendor storage environments and is a member of IBM Spectrum Storage™ Suite.
- **[IBM Virtual Storage Center](#)**
IBM Virtual Storage Center is a storage solution that provides efficient management, data protection, and virtualization for heterogeneous storage environments. It helps enhance storage efficiency, provides greater mobility, and delivers stronger control over storage performance and management.
- **[IBM Storage Insights for IBM Spectrum Control](#)**
IBM Storage Insights for IBM Spectrum Control is an IBM Cloud® service that can help you predict and prevent storage problems before they impact your business. It is complementary to IBM Spectrum Control and is available at no additional cost if you have an active license with a current subscription and support agreement for IBM Virtual Storage Center, IBM Spectrum Storage Suite, or any edition of IBM Spectrum Control.
- **[Replication products](#)**
IBM Spectrum Control no longer supports Tivoli® Storage Productivity Center for Replication. IBM Copy Services Manager is the replacement product for replication to use with IBM Spectrum Control.
- **[Product updates and security fixes](#)**
IBM Spectrum Control provides regular maintenance updates that can include code fixes, security fixes, new features, and enhancements.
- **[Key concepts](#)**
This section contains a technical overview that will help you understand how IBM Spectrum Control works. An understanding of the concepts in this section will help you use IBM Spectrum Control effectively.

Supported devices in IBM Spectrum Control

Find out which storage devices and products that you can monitor with IBM Spectrum® Control.

Storage systems

You can monitor the following storage systems with IBM Spectrum Control.

Table 1. Storage systems that can be monitored in IBM Spectrum Control

Storage System	Block	File	Object
Click a storage system to view its supported versions.			
DS8000®	✓		
Dell EMC Unity	✓	✓	
Dell EMC VMAX family	✓		
Dell EMC VNX, VNXe	✓	✓	
FlashSystem 5000	✓		
FlashSystem 5100	✓		
FlashSystem 5200	✓		
FlashSystem 7200	✓		
FlashSystem 7300	✓		
FlashSystem 9100	✓		
FlashSystem 9200	✓		
FlashSystem 9500	✓		
FlashSystem V9000	✓		
FlashSystem 900	✓		
FlashSystem A9000	✓		
FlashSystem A9000R	✓		
Hitachi VSP	✓		
IBM® Cloud Object Storage			✓
IBM Spectrum Accelerate	✓		
IBM Spectrum Scale (ESS and GSS)		✓	✓
IBM Spectrum Virtualize software-only clusters	✓		
IBM Spectrum Virtualize for Public Cloud	✓		
NetApp ONTAP 9		✓	✓
Pure FlashArray//M and FlashArray//X	✓		
SAN Volume Controller	✓		
Storwize® V3500	✓		
Storwize V3700	✓		
Storwize V5000	✓		
Storwize V7000	✓		
IBM Flex SystemFlashSystem V7000 Storage Node	✓		
Storwize V7000 Unified		✓	
XIV®	✓		
All others (managed by SMI-S providers)	✓		

Restriction: IBM Spectrum Control doesn't support monitoring non-IBM software-defined storage devices. However, it can monitor IBM software-defined storage devices, such as IBM SAN Volume Controller. For a list of storage devices that can be monitored, check out <https://www.ibm.com/support/pages/node/6249369>.

Tips:

- To monitor the performance of a Storwize V7000 Unified storage system, you must add it as a block storage system.
- To monitor an IBM Spectrum Virtualize for Public Cloud storage system, you must configure it for communication with IBM Spectrum Control. For more information, see [Configuring IBM Spectrum Virtualize for Public Cloud](#).

Switches and fabrics

You can monitor the following types of switches and fabrics with IBM Spectrum Control:

- Brocade
- Cisco

For a complete list of the supported switches, see <https://www.ibm.com/support/pages/node/6249365>.

Hypervisors

You can monitor the following hypervisors with IBM Spectrum Control:

- Virtual machines on KVM 2.0
- Virtual machines on VMware ESX, ESXi, and vCenter Server

For a complete list of supported hypervisors, see <https://www.ibm.com/support/pages/node/6249425#hypervisor>.

Servers

You can deploy Storage Resource agents to monitor the following servers with IBM Spectrum Control:

IBM AIX®

- AIX 7.3
- AIX 7.2
- AIX 7.1

Linux®

- Red Hat® Enterprise Linux 8
- Red Hat Enterprise Linux 7
- SuSE Linux Enterprise Server 11

Microsoft Windows

- Windows Server 2019
- Windows Server 2016
- Windows Server 2012 R2
- Windows Server 2012

For a complete list of supported servers, see <https://www.ibm.com/support/pages/node/6249425>.

Architecture

The IBM Spectrum® Control consists of several components that form the infrastructure of its storage-management functions.

Data server

This component is the control point for product scheduling functions, configuration, and event information. It also includes functions that schedule data collection and discovery for the Device server.

Device server

This component discovers, gathers information from, analyzes performance of, and controls storage subsystems and SAN fabrics. It coordinates communication with and data collection from agents that scan SAN fabrics and storage devices.

Web server

IBM Spectrum Control uses WebSphere® Application Server Liberty as the web application server to host its GUI, and storage management API for cloud.

Alert server

This component manages the complex event processing that is related to alerting on the condition of resources and their attributes. The ability to detect and be notified about configuration, capacity, and performance changes within a storage environment is important to helping you maintain and administer storage resources.

Export server

This component enables the export of data from the tables in the GUI to HTML format and the creation of custom reports that show information about capacity, configuration, and health status of the resources in your storage environment. It runs in the Node.js environment and you can produce custom reports for storage systems, back-end storage systems, and their internal resources such as disks, pools, and volumes.

IBM® Cognos® Analytics (optional)

If you plan to use the optional Cognos Analytics reports, you must install IBM Cognos Analytics. You can use Cognos Analytics to create reports and view predefined reports to analyze multiple storage systems, switches, servers, and hypervisors. You can install Cognos Analytics on the same computer with IBM Spectrum Control or you can install Cognos Analytics on a separate computer.

Database

A single database instance serves as the repository for all IBM Spectrum Control components.

Agents

Storage Resource agent, CIM agents, and SNMP agents gather host, application, storage system, and SAN fabric information and send that information to the Data server or Device server.

- GUI
Use the IBM Spectrum Control GUI to manage the storage infrastructure in an enterprise environment.
- CLI
Use the command-line interface (CLI) to issue commands for key IBM Spectrum Control functions.

Interfaces for IBM Spectrum Control

IBM Spectrum® Control provides multiple user interfaces for managing the storage infrastructure in an enterprise environment.

IBM Spectrum Control GUI

This interface runs in a web browser and includes the ability to monitor, manage, and troubleshoot storage resources. You can access this interface from anywhere that you have a web browser and connectivity to a network.

IBM® Cognos® Analytics (optional)

This interface runs in a web browser. Use this interface to view predefined reports and create custom reports about the storage systems managed by IBM Spectrum Control. For more information, see <https://www.ibm.com/docs/en/cognos-analytics/11.1.0>.

Command-line interface

Use this interface to run IBM Spectrum Control commands from a command prompt. For more information about CLI commands, see [Command-line interface](#).

Related concepts

- [Managing resources](#)

Starting IBM Spectrum Control

You can start IBM Spectrum® Control by opening a web browser and entering a web address for the IBM Spectrum Control logon page. For example, you might enter `https://storage.example.com:9569/srm`.

Before you begin

Before you start IBM Spectrum Control, ensure that you are using a supported web browser. For a list of web browsers that you can use with IBM Spectrum Control, see [IBM Spectrum Control - Platform Support: Servers, Agents, and Browsers - Web Browsers](#).

About this task

Start the IBM Spectrum Control GUI to administer and monitor the condition, capacity, and relationships of the resources within your storage environment.

Procedure

1. On a server running the Windows operating system, start IBM Spectrum Control GUI. For information about how to start the GUI on Windows operating systems, see [Opening IBM Spectrum Control GUIs and CLIs](#). If you are not on a server running the Windows operating system, start a web browser and enter the following address in the address field:

`https://host_name:port/srm`

In the preceding address, specify the following values:

host_name

The IBM Spectrum Control server. You can specify the host name as an IP address or a Domain Name System (DNS) name.

port

The port number for IBM Spectrum Control. The default port number for connecting to IBM Spectrum Control by using the HTTPS protocol is 9569. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation. If the default port number does not work, ask your IBM Spectrum Control administrator for the correct port number.


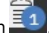

Tip: If you have a non-default port, check the value of the `WC_defaulthost_secure` property in `installation_dir/web/conf/portdef.props` file.

2. From the IBM Spectrum Control logon page, type your user name and password and click **Log in**.
The GUI opens in the browser.
Tip: If you want to log on to the GUI with Windows Domain credentials, use this form: `domain_name\user`.

Navigation

IBM Spectrum® Control provides many of the functions for managing a storage environment. To access the functions in its GUI, use the menu bar at the top of the main window.

GUI element	Description
-------------	-------------

GUI element	Description
Page banner	<p>Use the banner at the top of every page to complete the following tasks:</p> <ul style="list-style-type: none"> View the name of the product license that is active. Access resource pages and main product functions from the menu bar. View a list of failed tasks by using the failed tasks icon . You can retry or clear those tasks from the list. View a list of tasks that are running by using the running tasks icon . You can see how long the tasks are running. View the ID and role of the user who is logged in. The role of a user determines the product functions that are available to that user. For more information about roles, see Role-based authorization. Access the online help to view information about the currently displayed page and the overall product. Enter and exit rollup mode by using the rollup mode icon . You must add at least one subordinate rollup server for the rollup mode icon to display in the banner.
Home	<p>Dashboard</p> <p>Use the dashboard to view the overall status of monitored resources and identify potential problem areas in a storage environment. You can learn about:</p> <ul style="list-style-type: none"> Condition and usage of resources Entities that consume storage on those resources Number and status of unacknowledged alert conditions that are detected on the monitored resources Most active storage systems in your environment. <p>Performance Monitors</p> <p>Use performance monitors to collect information about the performance of storage systems and switches. This information includes key performance metrics and the percentage of data collections that succeeded during the most recent 24 hours when the performance monitor was active.</p> <p>Alerts</p> <p>Use alerts to be notified when certain configuration, status, and performance conditions are detected on monitored resources.</p> <p>Tasks</p> <p>Manage the tasks that are created when you complete the following actions:</p> <ul style="list-style-type: none"> Place volumes on the tiers that match the workload requirements of the volume. Move or convert volumes. Move volumes from overutilized pools to pools that are less used. <p>System Management</p> <p>View information about the overall condition of IBM Spectrum Control. You can view information about the servers on which the product is installed. This information includes component server and database status, certain server alerts and database connection alerts, server file-system capacity information, and remote volume-performance information.</p>
Storage	<p>Add storage systems for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer storage systems:</p> <ul style="list-style-type: none"> Modify the schedules for data collection. Place volumes on the tiers that match the workload requirements of the volume. Move or convert volumes. Move volumes from overutilized pools to pools that are less used. Add resources to a capacity pool. Open the management GUIs for storage systems. Update login credentials. Acknowledge conditions and statuses.
Servers	<p>Add servers and hypervisors for monitoring and view asset and status information about their internal and related resources. You can also complete the following actions to manage and administer these resources:</p> <ul style="list-style-type: none"> Deploy Storage Resource agents to servers for full monitoring. Modify the schedules for data collection. Modify and upgrade Storage Resource agents. Acknowledge conditions and statuses.
Network	<p>Add switches and fabrics for monitoring and view asset, status, and performance information about their internal and related resources. You can also complete the following actions to manage and administer these resources:</p> <ul style="list-style-type: none"> Modify the schedules for data collection. Setting or modifying the zoning policy. Open the management GUI for a switch. Acknowledge conditions and statuses.

GUI element	Description
Groups	<p>Add applications and departments to combine resources that are grouped by applications and departments. Applications and departments are used to monitor and view status, capacity, and performance information, and view details about their related resources. You can also complete the following actions to manage and administer applications and departments:</p> <ul style="list-style-type: none"> • Create applications and departments to model data for storage resources, allowing for enhanced capacity trending and performance troubleshooting. • View the status of resources that make up the applications and departments. • Create resource filters to automatically add resources to applications. • Add resources directly to applications and add applications to existing applications and departments. • Add departments to existing departments. • View subcomponent (member application) information to support hierarchal levels of storage grouping. • Remove applications and departments. <p>Review the distribution of capacity in your storage environment to determine whether you have sufficient capacity to meet the demands of your tiered storage.</p>
Advanced Analytics	<p>Configure IBM Spectrum Control for reclaiming storage that is not being used.</p> <p>Discontinued support: Other advanced analytics features, such as cloud configuration, provisioning, and optimization, are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see Discontinued features in IBM Spectrum Control.</p>
Reports	<p>Use the Reports option to create, configure, schedule, and send chargeback reports for applications, departments, hypervisors, and servers.</p> <p>Use IBM Spectrum Control when you want to individually view details about resources, for example when you want to troubleshoot a resource.</p> <p>Use the IBM® Cognos® Analytics reporting tool to analyze multiple storage systems, switches, servers, and hypervisors. For example, you can compare performance metrics across multiple storage systems. You can also view a report about the aggregate volume load that is on a server or hypervisor.</p> <p>Use the IBM Cognos Analytics reporting tool to schedule reports to run on different systems and with different options. You can specify schedules and output formats for the reports, and different ways to share the reports.</p> <p>For information about how to get started with the IBM Cognos Analytics reporting tool, see http://www.ibm.com/support/knowledgecenter/en/SSEP7J_11.0.0.</p>
Settings	<p>Complete the following actions to customize IBM Spectrum Control for your environment:</p> <ul style="list-style-type: none"> • Configure IBM Spectrum Control to send alert notifications by email, using SNMP traps, or using a Tivoli® Netcool®/OMNIBus server. An alert notification is sent when an alert condition is detected on the monitored resources in your environment. • Create and modify alert policies to manage the alert definitions and notification settings that apply to different sets of resources. • Modify the authentication repository and assign IBM Spectrum Control roles to user groups. Roles determine the product functions that are available to users. • Specify how long to retain the data that is collected about resources and the log files that are generated by IBM Spectrum Control. • Add and configure rollup servers so you can view capacity and status information about resources that are managed by your IBM Spectrum Control servers.

Keyboard navigation

Most of the features of the IBM Spectrum Control GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the GUI:

For navigating in the GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.
- To select an object, when the object is in focus, press Enter.

For actions menus:

- To navigate to the grid header, press Tab.
- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

For filters:

To specify a filter option and text:

1. Press Tab to navigate to the magnifying glass icon.
2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
3. Press Enter to select a filtering option.
4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

For text fields:

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the Submit button, press Tab.

For tables or lists:

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

Restriction: For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

Keyboard navigation with Firefox for Mac users: If you're using Firefox on a Mac with IBM Spectrum Control and want to use keyboard navigation, complete the following steps:

1. In Firefox, go to Preferences, > Advanced, > General and clear the check mark for Always use the cursor keys to navigate within pages. This step enables the use of Tab key to navigate between GUI elements.
2. In the URL address bar of Firefox, type about:config and press Enter.
Tip: If a warning prompt is displayed, click the button to accept the risk of changing browser settings. Existing settings won't be changed; instead, you'll be adding a preference setting for accessibility.
3. To add an accessibility preference for tab focus, right-click on the configuration page and select New, > Integer.
4. In the New integer value window, type accessibility.tabfocus and click OK.
5. Type 7 to set the integer value and click OK.
6. Open your Mac's System Preferences app, go to Keyboard, > Shortcuts, and select All Controls.

IBM Spectrum Storage Suite

IBM Spectrum Control provides monitoring, automation and analytics for multiple-vendor storage environments and is a member of IBM Spectrum Storage™ Suite.

Licensing options

IBM Spectrum Storage Suite gives you unlimited access to the IBM® Software Defined Storage product portfolio with licensing on a flat, cost-per-TB basis to make pricing easy to understand and predictable as storage capacity grows. You can save up to 40 percent compared with licensing the products separately. Non-production use of the software in test environments is included.

Learn more: Go to [IBM Spectrum Storage Suite at IBM Marketplace](#).

Which products are included with IBM Spectrum Storage Suite?

Besides the aforementioned IBM Spectrum Control, the following products are included in the IBM Spectrum Storage Suite:

Table 1. Products that are members of IBM Spectrum Storage Suite

Product	Key benefits	Links
IBM® Spectrum Protect™	Protects your data with multi-site replication and flexible restore capacity, and reduces your backup costs up to 53 percent. It can simplify data protection where data is hosted in physical, virtual, software-defined or cloud environments.	<ul style="list-style-type: none"> • Learn more and buy • Product documentation
IBM® Spectrum Protect™ Plus	Provides a data protection and availability solution for virtual environments that can be deployed in minutes and protect your environment within an hour. You can either implement it as a stand-alone solution, or integrate the solution with the IBM Spectrum Protect™ environment to offload copies for long-term storage and governance with scale and efficiency. IBM Spectrum Protect Plus on IBM Cloud® is a data protection and availability solution for virtual environments that can be deployed in minutes and protect your environment within an hour. IBM Spectrum Protect Plus on IBM Cloud can be provisioned with push-button simplicity to an existing or new VMware environment from the IBM Cloud console.	<ul style="list-style-type: none"> • Learn more and buy • Product documentation
IBM® Spectrum Scale™	Supports big data analytics and clustered applications with high-performance, scalable storage that enables global collaboration, simplifies workflows and lowers costs with cloud tiering.	<ul style="list-style-type: none"> • Learn more and buy • Product documentation
IBM® Spectrum Virtualize™	Provides an ideal way to manage and protect the huge volumes of data organizations use for big-data analytics and new cognitive workloads.	<ul style="list-style-type: none"> • Learn more and buy • Product documentation

Product	Key benefits	Links
IBM Spectrum Virtualize for Public Cloud	Provides a way to transition existing IT architecture to a hybrid-cloud or cloud-based model, where servers, storage, and network infrastructure are delivered in a public cloud environment.	<ul style="list-style-type: none"> Learn more and buy Product documentation
IBM® Spectrum Archive™	Provides a direct, intuitive and graphical access to data stored in IBM tape drives and libraries by incorporating the Linear Tape File System™ (LTFS) format standard for reading, writing and exchanging descriptive metadata on formatted tape cartridges.	<ul style="list-style-type: none"> Learn more and buy Product documentation
IBM® Spectrum Accelerate™	Provides a software-defined block storage solution that is designed for rapid and flexible deployment across heterogeneous infrastructure on and off premises. It can be deployed on your choice of servers: x86 commodity-choice hardware, integrated appliances (including IBM XIV® Storage System Gen3 and pre-validated third party appliances), and IBM public cloud.	<ul style="list-style-type: none"> Learn more and buy Product documentation
IBM® Cloud Object Storage	Provides highly scalable software-defined on-premises storage system designed to store, protect and easily access unstructured data. Deploy in your own data centers using IBM appliances or IBM software running on industry-standard, certified hardware.	<ul style="list-style-type: none"> Learn more and buy Product documentation
IBM Spectrum™ Discover	Modern metadata management software that provides data insight for petabyte-scale unstructured storage. The software easily connects to IBM Cloud Object Storage and IBM Spectrum Scale to rapidly ingest, consolidate and index metadata for billions of files and objects. It provides a rich metadata layer that enables storage administrators, data stewards and data scientists to efficiently manage, classify and gain insights from massive amounts of unstructured data.	<ul style="list-style-type: none"> Learn more and buy Product documentation

You might also be interested in the following products

Table 2. Other products of interest

Product	Key benefits	Links
IBM Storage Insights	An IBM Cloud based product that provides capabilities to help optimize your infrastructure with features, such as capacity planning, reclamation and data tier planning, and performance troubleshooting at a low monthly subscription price. Its enhanced dashboard enables you to quickly monitor and assess the basic health, status, and performance of block storage systems. If a problem is detected, you can get help to investigate and troubleshoot the problem and minimize the impact of hardware and software issues before they impact the performance of your critical business applications.	<ul style="list-style-type: none"> Learn more and buy Product documentation
IBM Spectrum™ Copy Data Management	Makes copies available to data consumers when and where they need them, without creating unnecessary copies or leaving unused copies on valuable storage. It catalogs copy data from across your local and hybrid cloud and off-site cloud infrastructure, identifies duplicates, and compares copy requests to existing copies.	<ul style="list-style-type: none"> Learn more and buy Product documentation

IBM Virtual Storage Center

IBM® Virtual Storage Center is a storage solution that provides efficient management, data protection, and virtualization for heterogeneous storage environments. It helps enhance storage efficiency, provides greater mobility, and delivers stronger control over storage performance and management.

Under a single license, IBM Virtual Storage Center includes the following products:

- IBM Spectrum® Control
- IBM Copy Services Manager
- IBM Spectrum Protect™ Snapshot
- IBM Spectrum Virtualize for IBM SAN Volume Controller, including all the functions available with storage virtualization, data reduction pools, HyperSwap® high availability, hybrid cloud and container support, and remote mirroring and FlashCopy® (snapshot).

The following key benefits are available with IBM Virtual Storage Center:

- Simplified storage management and reporting
- Virtualized physical storage resources for improved asset utilization
- Snapshot-based protection for the databases and applications that are used by IBM Spectrum Control
- Easy data mobility across pools and tiers of storage systems that use IBM Spectrum Virtualize
- Ability to change storage and move data without taking applications down
- Centralized management for visibility, control, and automation

- Helps improve the efficiency and productivity for storage management staff

For more details on what IBM Virtual Storage Center can do for your business, see <https://www.ibm.com/products/virtual-storage-center>

Licensing options

IBM Virtual Storage Center is available in three editions, each with different license options to best fit your needs:

Table 1. Licensing options for each of the IBM Virtual Storage Center editions

Product	Licensing	Why Choose This Edition
IBM Virtual Storage Center (perpetual license)	Measured by the number of storage capacity units (SCUs) that are needed to cover the total usable capacity in terabytes (TiB) of your storage.	<ul style="list-style-type: none"> • You use an IBM SAN Volume Controller as your storage virtualization engine and your total usable capacity is more than 1,500 TiBs. • You want to sign up for a perpetual license.
IBM Virtual Storage Center (subscription license)	Measured by the total usable capacity in terabytes (TiB) of your storage.	<ul style="list-style-type: none"> • You use an IBM SAN Volume Controller as your storage virtualization engine and your total usable capacity is more than 1,500 TiBs. • You want to pay for your license on a monthly basis with a flexible term length, have an initial lower cost, and use the familiar Passport Advantage® (PA) system to manage your contract.
IBM Virtual Storage Center Entry (formerly known as IBM Virtual Storage Center Entry Edition)	Measured by the total usable capacity in terabytes (TiB) of your storage, with the following limits: <ul style="list-style-type: none"> • Up to 1,500 TiBs of usable capacity • No limit I/O groups 	You use an IBM SAN Volume Controller as your storage virtualization engine and your total usable capacity is less than 1,500 TiBs.
IBM Virtual Storage Center for IBM Storwize®	Measured by the number of licenses that are needed for the storage enclosures, modules, or expansions of your storage.	You use an IBM Storwize or FlashSystem 5100 storage system in your environment.

When you subscribe to the IBM Virtual Storage Center license, you also receive IBM's Software Subscription and Support, which includes the following support options for the length of the subscription:

- Download access to the latest versions, releases, and fixes of the solutions in the IBM Virtual Storage Center license.
- IBM support for routine installation, deployment, migration, usage, and code-related technical support questions.
- 24x7 assistance for Severity 1 issues 7 days a week, 52 weeks a year. Simply open a case from the web, chat, or phone.

For the perpetual IBM Virtual Storage Center license, an SCU is the measure of capacity by which the license is charged. The number of SCUs is based on the drive classes that are used by your storage systems and the capacity that is assigned to each category. The following table shows the SCU categories, drive classes, and ratios:

Table 2. Supported SCU licenses, drive classes, and ratios

Category	Drive Classes	Ratio of usable storage to SCU
1	Storage Class Memory (SCM) drives and managed disks on IBM Spectrum Virtualize for Public Cloud	1 TiB to 1 SCU
2	Flash and Solid-State Drives (SSDs)	1 TiB to 0.847 SCU 1.18 TiB to 1 SCU
3	10K Serial Attached SCSI (SAS) Drives, 15K Fibre Channel Drives, and storage systems that use Category 4 drives with advanced architectures	1 TiB to 0.5 SCU 2 TiB to 1 SCU
4	Near Line SAS (NL-SAS), Serial ATA (SATA) Drives, and Safeguarded Copy (SGC) Target Capacity	1 TiB to 0.25 SCU 4 TiB to 1 SCU
Tips: <ul style="list-style-type: none"> • Any storage capacity that uses drive classes that are not listed in the table is classified as Category 1. • Calculations are rounded up to the nearest whole number. 		

What's included in IBM Virtual Storage Center

The following features are available in all editions of IBM Virtual Storage Center:

- Storage monitoring, reporting, and optimization for storage systems, hypervisors, servers, fabrics, and switches. Available on-premises or from IBM Cloud®.
- 3-site replication management
- Application-aware backup and restore operations with snapshot technologies of advanced storage systems
- External storage virtualization
- Easy Tier®
- FlashCopy
- Metro Mirror and Global Mirror

The following features are not included in any edition of IBM Virtual Storage Center:

- IBM Spectrum Virtualize Software for IBM SAN Volume Controller Real-time Compression Software V8
- IBM Spectrum Virtualize Software for IBM SAN Volume Controller Encryption Software

How to get IBM Virtual Storage Center

No matter which edition that you choose, IBM Virtual Storage Center can help save you time and money. It brings together storage management, monitoring, reporting, replication, and virtualization capabilities in one, comprehensive storage solution. You don't need to worry about shopping on your own for the best storage products to help meet your needs -- IBM has already done for that for you with IBM Virtual Storage Center.

To find out more about your purchasing options for IBM Virtual Storage Center, contact your IBM Business Partner or local IBM representative. In the US, to identify your IBM Business Partner or local IBM representative, you can call 800-IBM-4YOU (426-4968).

IBM Storage Insights for IBM Spectrum Control

IBM® Storage Insights for IBM Spectrum Control is an IBM Cloud® service that can help you predict and prevent storage problems before they impact your business. It is complementary to IBM Spectrum® Control and is available at no additional cost if you have an active license with a current subscription and support agreement for IBM Virtual Storage Center, IBM Spectrum Storage Suite, or any edition of IBM Spectrum Control.

As an on-premises application, IBM Spectrum Control doesn't send the metadata about monitored devices offsite, which is ideal for dark shops and sites that don't want to open ports to the cloud. However, if your organization allows for communication between its network and the cloud, you can use IBM Storage Insights for IBM Spectrum Control to transform your support experience for IBM block storage.

Why use IBM Storage Insights for IBM Spectrum Control

IBM Storage Insights for IBM Spectrum Control and IBM Spectrum Control work hand in hand to monitor your storage environment. Here's how IBM Storage Insights for IBM Spectrum Control can transform your monitoring and support experience:

- Open, update, and track IBM Support tickets easily for your IBM block storage devices.
- Get hassle-free log collection by allowing IBM Support to collect diagnostic packages for devices so you don't have to.
- Use Call Home to monitor devices, get best practice recommendations, and filter events to quickly isolate trouble spots.
- Leverage IBM Support's ability to view the current and historical performance of your storage systems and help reduce the time-to-resolution of problems.

To compare the features of **IBM Spectrum Control** and **IBM Storage Insights for IBM Spectrum Control**, check out the [Feature comparison](#).

Important: To monitor the devices that you already monitor in IBM Spectrum Control, you must add them separately to IBM Storage Insights for IBM Spectrum Control. You must also deploy a data collector in IBM Storage Insights for IBM Spectrum Control to collect advanced metadata about devices.

Alerting tip: Alerts are a good way to be notified of conditions and potential problems that are detected on your storage. If you use **IBM Spectrum Control** and **IBM Storage Insights for IBM Spectrum Control** together to enhance your monitoring capabilities, it's recommended that you define alerts in one of the offerings and not both. By defining all your alerts in one offering, you can avoid receiving duplicate or conflicting notifications when alert conditions are detected.

How to get IBM Storage Insights for IBM Spectrum Control

To register for IBM Storage Insights for IBM Spectrum Control, go to <https://www.ibm.com/it-infrastructure/storage/storage-insights/registration> and follow the directions for current IBM Spectrum Control customers.

To see how easy it is to sign up and deploy a data collector, check out [Before you begin checklist for IBM Storage Insights](#).

Questions and answers

Find answers to other questions about IBM Storage Insights for IBM Spectrum Control.

What's the difference between IBM Storage Insights, IBM Storage Insights Pro, and IBM Storage Insights for IBM Spectrum Control?

- **IBM Storage Insights** is an off-premises, IBM Cloud service that is available free of charge if you own IBM block storage systems. It provides a unified dashboard for IBM block storage systems with a diagnostic events feed, a streamlined support experience, and key capacity and performance information.
- **IBM Storage Insights Pro** is an off-premises, IBM Cloud service that is available on subscription and expands the capabilities of IBM Storage Insights. You can monitor IBM file, object, and software-defined storage (SDS) systems, and non-IBM block and file storage systems such as Dell EMC storage systems. It also includes configurable alerts and predictive analytics that help you to reduce costs, plan capacity, and detect and investigate performance issues. You get recommendations for reclaiming unused storage, recommendations for optimizing the placement of tiered data, capacity planning analytics, and performance troubleshooting tools.
- **IBM Storage Insights for IBM Spectrum Control** is similar to IBM Storage Insights Pro in capability and is available for no additional cost if you have an active license with a current subscription and support agreement for IBM Virtual Storage Center, IBM Spectrum Storage Suite, or any edition of IBM Spectrum Control.

I already have a paid subscription to IBM Storage Insights Pro, does this service apply to me?

Yes! When your next renewal date for IBM Storage Insights Pro approaches, contact your IBM Sales representative to see if a move to IBM Storage Insights for IBM Spectrum Control is right for you.

Tip: If you want to cancel your subscription to IBM Storage Insights Pro and switch to IBM Storage Insights for IBM Spectrum Control, check out the IBM Cloud Service Agreement at <https://www.ibm.com/support/customer/csol/contractexplorer/cloud/csa/us-en/9> for cancellation information.

If I already have IBM Storage Insights (the free service), can I upgrade it to IBM Storage Insights for IBM Spectrum Control?

Yes, if you have an active license for IBM Spectrum Control, you can upgrade IBM Storage Insights to IBM Storage Insights for IBM Spectrum Control. Just go to <https://www.ibm.com/it-infrastructure/storage/storage-insights/registration>, choose the option for IBM Spectrum Control, and follow the prompts.

Because IBM Storage Insights for IBM Spectrum Control is similar to IBM Storage Insights Pro in capability, does it use the same Service Level Agreement (SLA)?

No, IBM Storage Insights for IBM Spectrum Control doesn't include the service level agreement for IBM Storage Insights Pro. Terms and conditions for IBM Storage Insights for IBM Spectrum Control are available at <http://www.ibm.com/software/sla/sladb.nsf/sla/sd-8410-01>.

Do IBM Spectrum Control and IBM Storage Insights for IBM Spectrum Control share the metadata that they collect about storage?

No, metadata cannot be shared between the offerings or exported from one offering to the other.

Are the devices that I monitor in IBM Spectrum Control automatically monitored by IBM Storage Insights for IBM Spectrum Control?

To monitor the devices that you already monitor in IBM Spectrum Control, you must also add them to IBM Storage Insights for IBM Spectrum Control.

To collect advanced metadata about devices, you must also deploy a data collector in IBM Storage Insights for IBM Spectrum Control. Storage Resource agents and device connections that are defined in IBM Spectrum Control aren't used in IBM Storage Insights for IBM Spectrum Control.

To see which devices that you can monitor in IBM Storage Insights for IBM Spectrum Control, check out the list for IBM Storage Insights Pro in [Supported resources](#).

If I use IBM Storage Insights Pro or IBM Storage Insights and switch to IBM Storage Insights for IBM Spectrum Control, is my existing metadata migrated?

Yes, the metadata that you collected in IBM Storage Insights Pro and IBM Storage Insights is automatically available in IBM Storage Insights for IBM Spectrum Control.

How is licensing handled?

IBM Storage Insights for IBM Spectrum Control uses the same licensing that you already have in place for IBM Spectrum Control. IBM Spectrum Control can be licensed by capacity, by storage capacity units (SCU), or by number of slots in enclosures.

How long can I use IBM Storage Insights for IBM Spectrum Control?

You can use IBM Storage Insights for IBM Spectrum Control for as long as you have an active license with a current subscription and support agreement for IBM Spectrum Control license. If your subscription and support lapses, you're no longer eligible for IBM Storage Insights for IBM Spectrum Control.

If your subscription and support lapses, don't worry. To continue using IBM Storage Insights for IBM Spectrum Control, simply renew your IBM Spectrum Control license.

You can also choose to subscribe to IBM Storage Insights Pro. For information about how to subscribe, see [Want to try or buy IBM Storage Insights Pro?](#)

Because IBM Storage Insights for IBM Spectrum Control is a cloud service, I'm concerned about security. Where can I find out more information about its security measures?

IBM Storage Insights for IBM Spectrum Control runs in IBM Cloud and adheres to IBM's rigorous security standards. For more information, check out the [\[PDF\] IBM Storage Insights Security Guide](#).

Where can I learn more about IBM Storage Insights?

Explore these links for more detailed information about IBM Storage Insights:

- [IBM Storage Insights Knowledge documentation](#)
- [IBM Storage Insights FAQ](#)

How do I get support for IBM Storage Insights for IBM Spectrum Control?

To contact IBM Support for help and report issues that you encounter in IBM Storage Insights for IBM Spectrum Control, follow these steps:

1. [Open a support case](#) against IBM Storage Insights.
2. Describe the problem. To help us troubleshoot, include the URL of your IBM Storage Insights for IBM Spectrum Control instance.
3. Submit the case.

Feature comparison

Compare the features of IBM Spectrum Control and IBM Storage Insights for IBM Spectrum Control.

Table 1. Features in IBM Spectrum Control and IBM Storage Insights for IBM Spectrum Control

Resource Management	Features	IBM Spectrum Control	IBM Storage Insights for IBM Spectrum Control
Monitoring	Inventory	IBM and non-IBM block storage, file storage, object storage, hypervisors, fabrics, and switches	
	Call Home events		✓
	Performance	✓ (1-minute intervals)	✓ (5-minute intervals)
	Capacity	✓	✓
	Drill down performance workflows to troubleshoot bottlenecks	✓	✓
	Explore virtualization relationships	Storage and Server virtualization	Storage virtualization
	Explore replication relationships	✓	✓
	Retain performance data	Customizable	1 year
Service	Deployment method	On-premises	Off-premises (in IBM Cloud)
	Filter Call Home events to quickly isolate trouble spots		✓
	Hassle-free log collection		✓
	Simplified ticketing		✓
	Show active PMRs and ticket history		✓
	Active directory and LDAP integration for managing users	✓	
Reporting	Inventory, capacity, performance, and storage consumption reports	✓	✓
	Rollup reporting	✓	
	REST API	✓	
Alerting	Predictive Alerts	✓	✓
	Customizable, multi-conditional alerting, including alert policies	✓	✓
Analytics	Performance planning	✓	✓
	Capacity planning	✓	✓
	Business impact analysis (applications, departments, and groups)	✓	✓
	Provisioning with service classes and capacity pools	✓	
	Balance workload across pools	✓	
	Optimize data placement with tiering	✓	✓
	Optimize capacity with reclamation	✓	✓
	Transform and convert volumes	✓	
Pricing		On-premises licensing	No charge for IBM Spectrum Control customers

Replication products

IBM Spectrum® Control no longer supports Tivoli® Storage Productivity Center for Replication. IBM® Copy Services Manager is the replacement product for replication to use with IBM Spectrum Control.

Information about Copy Services Manager

It is strongly recommended that you back up your replication environment before you proceed with any upgrade.

See "Backing up a replication environment" and verify that your replication backup file location is outside of the IBM Spectrum Control installation directory structure or the replication backup file is deleted during the uninstall of replication.

Copy Services Manager provides automation for 2-site replication, 3-site replication, and advanced copy services from a single point of control across multiple platforms. To learn more about how to download and install Copy Services Manager, see [IBM Copy Services Manager](#).

Downloading the software code and the license for Copy Services Manager

Before you migrate to Copy Services Manager, obtain the software code and the license.

To download the software code for Copy Services Manager, go to IBM Fix Central at <https://www.ibm.com/support/fixcentral>.

To download the license for Copy Services Manager, follow these steps to access the IBM Passport Advantage® Online download portal. The license is part of the IBM Spectrum Control eAssembly in IBM Passport Advantage.

Note: You need an IBM ID and password to access Passport Advantage Online for customers.

1. Click [Downloading IBM Spectrum Control](#).
2. Scroll to **Download Package**.
3. Click **Downloading IBM Spectrum Control Vvrm using Passport Advantage Online** (where *vrm*=the version, release, and modification of the software).
4. Scroll to the **Optional parts**, and locate the part number for Copy Services Manager. Make a note of the part number.
5. Enter <https://www.ibm.com/software/passportadvantage/pacustomers.html> to go to Passport Advantage Online for customers.
6. Click **Customer sign in** to enter your IBM ID and password. If you are not an authorized user, follow the steps on the page to request access.
7. On the **Software and services online** page, click **Software download & media access**.
8. Use the **Find by part number** search option, and enter the part number that you copied in Step 4 to download the license.

Migrating to Copy Services Manager

Before you migrate to Copy Services Manager, be sure to back up your configuration of Tivoli Storage Productivity Center for Replication. Then, install Copy Services Manager and uninstall Tivoli Storage Productivity Center for Replication.

To perform the migration, follow these steps:

1. Back up your Tivoli Storage Productivity Center for Replication configuration by using the steps in the topic "Backing up a replication environment." The topic is located in the IBM Spectrum Control 5.3.3 Knowledge documentation or in the IBM Spectrum Control 5.3.3 Installation Guide. You can find and download the guide from the "Printable documentation" topic in the IBM Spectrum Control 5.3.3 Knowledge documentation.
2. Download the software code for Copy Services Manager (go to IBM Fix Central at <https://www.ibm.com/support/fixcentral> and follow the steps).
3. Download the license for Copy Services Manager (see the preceding section **Downloading the software code and the license for Copy Services Manager**).
4. Install Copy Services Manager. For instructions, see [Installing on distributed systems](#).
Note: During the installation, you will be asked to choose whether to "Migrate or restore by using existing backup." For details, see [Migrating on distributed systems](#).
5. Uninstall Tivoli Storage Productivity Center for Replication by using the steps in the topic "Manually uninstalling Tivoli Storage Productivity Center for Replication." The topic is located in the IBM Spectrum Control 5.2.12 Knowledge documentation or in the IBM Spectrum Control 5.2.12 Installation Guide. You can find and download the guide in the "Printable documentation" topic in the IBM Spectrum Control 5.2.12 Knowledge documentation.

Product updates and security fixes

IBM Spectrum® Control provides regular maintenance updates that can include code fixes, security fixes, new features, and enhancements.

- [Releases and downloads](#)
IBM Spectrum Control provides regular maintenance updates that can include code fixes, security fixes, new features, and enhancements.
- [Security](#)
In IBM Spectrum Control, the security of your data is critical. At IBM®, we're serious about keeping that data safe and constantly strive to address vulnerabilities before they impact your business.
- [Subscribing to IBM announcements](#)
Subscribe to My Notifications to be notified automatically of IBM announcements such as security bulletins and flashes for IBM Spectrum Control.
- [End of support](#)
End of Support (EOS) marks the official withdrawal of technical support for specific versions and releases of IBM Spectrum Control and other IBM software products.
- [Collaborating with the team](#)
Collaborate with the IBM Spectrum Control team to help improve the product.
- [Reporting a problem](#)
Report problems that you encounter in IBM Spectrum Control.

Releases and downloads

IBM Spectrum® Control provides regular maintenance updates that can include code fixes, security fixes, new features, and enhancements.

Three types of maintenance packages are available for IBM Spectrum Control:

- Releases represent a significant update to the product, which includes, but is not limited to, new features, enhancements to existing features, code fixes, and security fixes.
- Refresh packs are quarterly deliverables that update the modification level of a release, and include new features, enhancements to existing features, code fixes, and security fixes.
- Fix packs update the fix level of a release, and primarily include code and security fixes.

All types of maintenance packages are cumulative: a fix pack includes all previous fix packs; a refresh pack includes all the previous fix packs and refresh packs; a release includes all previous fix packs, refresh packs, and releases.

With each new maintenance package that is delivered, the product number changes. The following table provides example of product number changes when you install a package.

Maintenance packages	Product version change examples
Release	5.4
Refresh pack (modification level)	5.4.1
Fix pack	5.4.1.1

A fix pack is installed on top of a specific modification level and release. If you install a fix pack that belongs to a newer level than your current fix pack, then your environment is automatically upgraded to the newer fix pack. For example,:

- If you are on fix pack 1 (5.4.1.1) and you apply fix pack 2, then the version is changed to 5.4.1.2.
- If you are on fix pack 1 (5.4.1.1) and you apply fix pack 3, then your environment is automatically updated to fix pack 3 and your version is changed to 5.4.1.3.

Tip: When you install a release, the modification level and fix pack number are reset to 0.

Downloading IBM Spectrum Control

If you have an IBM® ID and have purchased IBM Spectrum Control, you can download packages for a new installation or download packages that you can apply as an upgrade or fix to your existing installation.

- To download the package for a new installation of IBM Spectrum Control, go to [Passport Advantage Online](#).
- To download the package for upgrading or applying the latest fixes to your existing installation of IBM Spectrum Control, go to [Fix Central](#).

Tips:

- If you are downloading the package for a new installation, a license is required and is included on Passport Advantage® Online.
- If you are downloading the package for an upgrade, your existing license is transferred.
- Fix Central is the recommended site for retrieving fixes for most IBM products. Fix Central enables you to search, select, and download appropriate fixes for their situation.

Security

In IBM Spectrum® Control, the security of your data is critical. At IBM®, we're serious about keeping that data safe and constantly strive to address vulnerabilities before they impact your business.

Maintenance and security fixes are regularly provided for the IBM Spectrum Control base product. The base product includes core technology and components such as WebSphere® Application Server Liberty profile and the Java™ Runtime Environment.

When security issues or vulnerabilities are discovered in IBM Spectrum Control or its core components, IBM analyzes the product and distributes a fix or mitigation instructions, as appropriate. These issues are communicated through security bulletins on ibm.com® and through discrete communication when necessary.

To find security bulletins, go to the following pages on ibm.com:

- [Security bulletins for IBM Spectrum Control](#)
- [IBM Support portal](#)
- [Supported storage products](#)

More information about security: For more information about security in IBM Spectrum Control, check out this topics:

- [List of fixes that were included in releases of IBM Spectrum Control](#)
- [Ports used by IBM Spectrum Control](#)
- [Required user roles for monitoring resources](#)
- [User names and passwords](#)
- [Planning for IBM Spectrum Control authentication and authorization](#)
- [Planning to use LDAP for IBM Spectrum Control authentication](#)

Managing security for related software

An IBM Spectrum Control installation also includes other products, such as Db2®. You can monitor fixes for these related products in the same way that you do for IBM Spectrum Control on ibm.com:

- IBM support portal for Db2: https://www.ibm.com/support/entry/myportal/product/information_management/db2_for_linux_unix_and_windows
- IBM support portal for IBM WebSphere Application Server: https://www.ibm.com/support/entry/myportal/product/websphere/websphere_application_server
- IBM Product Security Response Team blog at <https://www.ibm.com/blogs/psirt/>

Each product team is responsible for publishing their own security bulletins. Review any published security bulletins to identify affected versions and follow the remediation or fix information. If these products are affected, they might provide a maintenance fix. It is acceptable to apply a fix pack to these products to resolve security vulnerabilities or other problems.

Important: Upgrading to a new version or release of these products is not typically supported by IBM Spectrum Control and might result in compatibility issues. Even with fix pack updates, read the security bulletins carefully for any compatibility concerns. For example, if a common protocol is disabled, both the related product and IBM

Spectrum Control must also accommodate a new protocol.

Subscribing to IBM announcements

Subscribe to My Notifications to be notified automatically of IBM® announcements such as security bulletins and flashes for IBM Spectrum® Control.

About this task

With My Notifications, you can specify that you want to receive daily or weekly email announcements. You can specify what type of information you want to receive (such as publications, hints and tips, product flashes (also known as alerts), downloads, and drivers). You can also customize and categorize the products about which you want to be informed and the delivery methods that best suit your needs.

Procedure

1. Go to <http://www.ibm.com/software/support/einfo.html>.
2. Click Subscribe now!.
3. Log in with your IBMid.
4. In Product lookup, type IBM Spectrum Control.
5. Click Subscribe for the edition of IBM Spectrum Control that you use.
6. Select the types of documents for which you want to receive notifications and click Submit.

End of support

End of Support (EOS) marks the official withdrawal of technical support for specific versions and releases of IBM Spectrum® Control and other IBM® software products.

What does End of Support mean to you

Important:

The following releases have reached end of support:

- Tivoli® Storage Productivity Center 5.2.0 - 5.2.8, 5.2.8 - 5.2.17
- IBM Spectrum Control 5.2.8 - 5.2.17
- IBM Spectrum Control 5.3.0 - 5.3.7

When a specific version or release of IBM Spectrum Control reaches end of support (sometimes called end of life), you might find that access to entitled severity 1 technical support is no longer available. The documentation for each product version is still available in PDF format when you access [IBM Documentation](#).

What can you do to monitor End of Support dates

Review the End of Support Announcement letters regularly, see <https://www.ibm.com/software/passportadvantage/endofsupportannouncementletters.html>

End of support announcements are generally made twice a year in April and September and are subject to change. You can confirm dates by visiting the IBM Software Support Lifecycle policy site, see <https://www.ibm.com/support/home/pages/lifecycle/index.html>

For a complete list of supported versions of IBM Spectrum Control, see <https://www.ibm.com/support/pages/overview-ibm-spectrum-control-releases>

End of support options

The best option for you when a version or release of IBM Spectrum Control is going into end of support is to upgrade to a newer version or release. You not only maintain access to new fixes and features, but you can open support cases that ensure if you do have an issue that requires immediate attention that you can get it.

You can purchase an IBM Software Support Service Extension (also known as Extended Support). Service Extensions give you access to product fixes and the ability to engage with IBM Support for use and defect support after your End of Support date is reached. Note, these extensions are not available for all products. Ask your IBM Support representative if IBM Spectrum Control is eligible for the extension.

Collaborating with the team

Collaborate with the IBM Spectrum® Control team to help improve the product.

About this task

Got a great idea for making IBM Spectrum Control even better? Do you want to vote for an enhancement that was requested by another user? The IBM® System Storage® Ideas Portal is a place where you can collaborate with the development team for IBM Spectrum Control and other users through your ability to search, view, comment on, submit and track product requests. Help shape the future of IBM!

Procedure

1. Go to the IBM System Storage Ideas Portal at <https://ibm-sys-storage.ideas.ibm.com/ideas?project=SPCO>.

2. To submit an enhancement request, click Add A New Idea.
3. Log in with your IBMid.
4. In the Choose a workspace for this idea drop-down menu, select IBM Spectrum Control .
5. Complete the form and click Add Idea.

Reporting a problem

Report problems that you encounter in IBM Spectrum® Control.

Before you begin

IBM® is committed to customer satisfaction and is ready to assist if you encounter any problems with IBM Spectrum Control. By opening a case, you can ensure that IBM Support will be informed of your problem. When you open a case, have the following information ready:

- The version, release, modification, and service level number of your IBM Spectrum Control installation.
- The communication protocol (for example, TCP/IP), version, and release number that you are using.
- The activity that you were doing when the problem occurred, listing the steps that you followed before the problem occurred.
- The exact text of any error messages.

Procedure

To open a case for IBM Spectrum Control, complete the following steps:

1. Go to the [IBM Support portal](#).
2. Click Open a case.
3. Log in with your IBMid.
4. Describe the problem.
5. Submit the case.

Videos: The IBM Support portal is supported by IBM Watson® and provides you with enhanced transparency into ticket resolution workflow and improved self-service options. Watch a few short videos to learn more:

- [Introducing A New Customer Portal](#)
- [Open And Manage Cases](#)
- [IBM Support Community: Search](#)
- [IBM Support Community: Forums](#)

Key concepts

This section contains a technical overview that will help you understand how IBM Spectrum® Control works. An understanding of the concepts in this section will help you use IBM Spectrum Control effectively.

- **[Data collection](#)**
To help you implement a storage management strategy, it is critical that you determine the information that you want to gather about the resources within your environment. You can schedule different data collection jobs depending on the types of information that you want to gather and the resources that you want to monitor.
- **[Performance monitoring and troubleshooting](#)**
IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes key performance metrics and alerts of threshold violations that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.
- **[Applications, departments, and general groups](#)**
Use applications and departments to organize the resources in your storage environment into a hierarchy that matches the structure of your business organization. The information can be used to view the performance of resources that belong to an application and the capacity growth of a department. Use general groups to group your resources, such as the storage systems with lease agreements that end in the current year, so that you can view information about the resources at one location in the GUI.
- **[Storage optimization](#)**
To optimize the resources in your storage environment, you can create tasks to balance pools, re-tier volumes, or transform volumes that are in storage virtualizer pools.
- **[Cloud configuration](#)**
To take advantage of the simplified provisioning and optimization capabilities in IBM Spectrum Control, some configuration is required.
- **[Provisioning storage overview](#)**
IBM Spectrum Control guides you through the steps for provisioning storage volumes or network-attached storage (NAS) shares. You can provision storage volumes or NAS file shares to one or more servers, one or more hypervisors, or one cluster.
- **[Storage reclamation](#)**
You can use volume reclamation recommendations to reclaim unused storage capacity in your environment and use your storage more efficiently.
- **[Alerts and alert policies](#)**
Specify conditions that trigger alerts and the actions to take when those alerts are triggered, such as notify an email address. Use alert policies to define those alert conditions and notification settings for a group of resources.
- **[Reporting](#)**
Use reporting functions to view overview and detailed information about your storage.
- **[Rollup servers](#)**
Rollup reporting combines capacity and status information from multiple instances of IBM Spectrum Control for monitoring and reporting. In rollup mode, secondary rollup servers collect information from managed resources and roll that data up to primary rollup servers.
- **[Units of measurement for storage data](#)**
IBM Spectrum Control uses decimal and binary units of measurement to express the size of storage data.
- **[Role-based authorization](#)**
Roles determine the functions that are available to users of IBM Spectrum Control. When a user ID is authenticated to IBM Spectrum Control through the GUI, CLI,

or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

- **[Fabrics and zones](#)**

Use the IBM Spectrum Control to learn more about fabrics and zones.

- **[Agents](#)**

The IBM Spectrum Control uses agents to gather data: Common Information Model (CIM) agents, Storage Resource agents, and SNMP agents.

- **[IBM Spectrum Control REST API](#)**

You can use the Representational State Transfer (REST) API for IBM Spectrum Control to access information about resources and to generate custom capacity, configuration, and performance reports.

Data collection

To help you implement a storage management strategy, it is critical that you determine the information that you want to gather about the resources within your environment. You can schedule different data collection jobs depending on the types of information that you want to gather and the resources that you want to monitor.

Use probes and performance monitors to collect detailed information about resources. You must schedule and run these data collection jobs before you can complete other tasks within IBM Spectrum® Control, such as viewing reports, setting alerts, and managing storage systems.

You schedule probe and performance monitor jobs for a resource when you add the resource for monitoring. After you add the resource, you can modify the probe and performance monitor job schedules on the list and details pages for the resource. For example, to modify the probe job schedule for a storage system, go to the Storage Systems page or the Storage System details page in the GUI.

Schedule data collection jobs to gather the following types of information:

Asset and status information

Use [probes](#) to collect asset, storage statistics, and status information about resources. You can run probes on all the resources that are monitored by IBM Spectrum Control.

Performance monitoring

Use [performance monitors](#) to collect metrics that measure the performance of switches and storage systems.

Performance monitoring and troubleshooting

IBM Spectrum® Control can collect information about the performance of storage systems and switches. This information includes key performance metrics and alerts of threshold violations that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.

To monitor the performance of resources and check for threshold violations, complete the following tasks:

- Add resources for monitoring and schedule data collection
- Define alerts for performance thresholds
- View and troubleshoot performance issues

Collect performance data

Before you can troubleshoot and view reports about performance, you must collect data about monitored resources. Performance monitors are data collection jobs that gather performance information about resources. This information includes metrics that measure the performance of the components within a resource. Metrics measure the performance characteristics of volumes, ports, and disks on storage systems and switches. IBM Spectrum Control provides many different metrics for measuring performance. For example, some key metrics for storage systems are I/O rate in I/O operations per second, data rate in MiB per second, and response time in milliseconds.

You can use metrics in IBM Spectrum Control to track growth or change in I/O rates, data rates, and response times. In many environments, I/O and data rates grow over time, and response times increase as those rates increase. This relationship can help with "capacity planning" for your storage. As rates and response times increase, you can use these trends to project when more storage performance and capacity is required.

Define alerts for performance thresholds

Alerts can notify you when the performance of a monitored resource falls outside of a specified range and might represent a potential problem. When you define an alert for an internal component of a resource, select a specific metric that you want to measure and its boundary values. When the performance of a resource falls outside the boundary values, an alert is triggered.

For example, you can define an alert that is triggered when the overall back-end response time for a managed disk on a SAN Volume Controller exceeds a certain value. The overall back-end response time is a metric that measures the average number of milliseconds that it takes to service each I/O operation on a managed disk.

View and troubleshoot performance issues

After data collection and performance thresholds are configured, you can use the web-based GUI to complete the following tasks:

- Measure, compare, and troubleshoot the performance of switches, storage systems, and their internal resources.
- Review the threshold violations and alerts that were triggered when the performance of a resource fell outside of a specific range.
- View performance information in a chart or table format to help you quickly identify where and when performance issues are occurring. The chart is a visual representation of how the performance of resources trend over time.
- Customize views of performance so that you can analyze specific resources and metrics during time ranges that you specify.
- Drill down into resources to view detailed information about the performance of internal and related resources. For example, if a SAN Volume Controller storage system is shown in the chart, you can quickly view and compare the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server for monitoring, IBM Spectrum Control automatically correlates that server with the ports on known host connections. If matches are found between the server and host connections on

monitored storage systems, you can view the performance of the internal resources that are directly associated with the SAN storage that is assigned to the server. For example, if a SAN Volume Controller maps two volumes to the server, you can view the performance of those volumes and the related managed disks.

- Export performance information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.
- In the optional Cognos® Analytics reporting tool, you can also view and create performance reports about multiple resources.

Related concepts

- [Reporting with Cognos Analytics](#)
- [Monitoring the performance of resources](#)

Related tasks

- [Adding resources](#)

Related reference

- [Performance metrics](#)

Applications, departments, and general groups

Use applications and departments to organize the resources in your storage environment into a hierarchy that matches the structure of your business organization. The information can be used to view the performance of resources that belong to an application and the capacity growth of a department. Use general groups to group your resources, such as the storage systems with lease agreements that end in the current year, so that you can view information about the resources at one location in the GUI.

IBM Spectrum® Control provides storage grouping to define resources that utilize application and department business modeling. The information can be used to view the performance of resources that belong to an application and the capacity growth of a department. The pages for the monitoring and management of applications and departments are aligned with user scenarios, for example, capacity trending, resource health monitoring and performance troubleshooting

To manage and administer applications and departments, complete the following tasks:

- Create applications and departments to model data for storage resources, allowing for enhanced capacity trending and performance troubleshooting.
- View the status of resources that make up the applications and departments.
- Create resource filters to automatically add resources to applications.
- Add resources directly to applications.
- Add applications to existing applications and departments.
- Add departments to existing departments.
- View subcomponent (member application) information to support hierarchical levels of storage grouping.

Applications

An application is a program or a project that consumes storage resources within an organization and interfaces with other enterprise groups that are important to the running of a business. Use IBM Spectrum Control to model the storage usage that is consumed in your environment by assigning the usage to applications to see the overall health status.

An application can be part of a department and have its own subcomponents (member applications) that are used to create a five level deep hierarchy. Applications that are grouped together can range from large line-of-business systems to specialized software, in a department, that runs on either client computers or servers. For example, an application might be an automated billing system within the Finance department, VMware running in the Information Technology department or an email marketing system that is part of the Marketing department.

The Applications page shows capacity information about the applications that are monitored by organization and interfaces with other enterprise groups that are important to the running of a business. Use IBM Spectrum Control to view and manage the resources that are defined to the applications. If the application is associated with a department, information about that department is also shown.

Departments

A department is a division within a business. Use IBM Spectrum Control to model the storage capacity that is consumed in your department for your business environment, in accordance with other department members.

A department can be hierarchical in its organizational layout. For example, a department might use 15 applications and be part of another five departments. A department might share storage resources with another department, subdepartment or an application even if they don't belong in the same hierarchy. For example, in a collaboration scenario, a single IBM® SAN Volume Controller might be shared by multiple departments.

The Departments page shows storage capacity information about the top-level departments, the subdepartments and any applications that belong to the department that are monitored by IBM Spectrum Control.

General groups

Create general groups to quickly view information about storage resources that have common characteristics. For example, you might group the subset of ports on a SAN Volume Controller that are used for inter-node communication or the storage systems that are used by a critical business application.

You can organize your storage resources into a general group hierarchy. Organizing resources into general groups and their subgroups can be helpful when you want to quickly view information about a group of resources, but you also want to view information about subgroups of resources within the group. For example, you can group the

resources that are used by your production application so that you can monitor all the application resources and separately monitor the specific subgroups of storage systems and ports.

Storage optimization

To optimize the resources in your storage environment, you can create tasks to balance pools, re-tier volumes, or transform volumes that are in storage virtualizer pools.

Discontinued support: Storage optimization is no longer supported in IBM Spectrum® Control. While the feature might still work in this release, it's recommended that you use another tool for your optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

Balancing the workload of volumes across storage pools

You can balance the workload of volumes across pools on the same tier. The pools are analyzed and recommendations are generated to move volumes from pools with high-activity values to pools with low-activity values.

Re-tiering volumes

You can re-tier volumes to balance pools or to tier volumes based on the criteria that you set in tiering policies. For example, you can tier volumes that are based on the volume workload or on file usage, or both. Depending on the conditions that are set in the tiering policy, recommendations are generated. For example, you can reduce storage costs by moving volumes with low workloads to lower or less expensive tiers. You can also improve performance and use storage more efficiently by moving volumes with heavy workloads to the tiers that best meet their workload requirements.

Transforming volumes

You can complete the following tasks for one or more volumes in storage virtualizer pools:

- Move volumes from one storage virtualizer pool to another pool on the same storage virtualizer.
- Move volumes in a storage virtualizer pool to a pool that is enabled for Easy Tier®.
- Convert fully allocated volumes to thin-provisioned volumes and convert thin-provisioned volumes to fully allocated volumes.
- Convert fully allocated volumes to compressed volumes and convert compressed volumes to fully allocated volumes.

The pools are analyzed and recommendations are generated. For example, analysis occurs to ensure that there is sufficient space in the pool to convert volumes or add volumes.

Related tasks

- [Optimizing storage pools](#)
- [Optimizing storage tiering](#)
- [Transforming and migrating volumes](#)
- [Managing tasks for tiering storage, balancing pools, and transforming storage](#)

Related reference

- [Criteria for identifying the pools that require balancing](#)

Cloud configuration

To take advantage of the simplified provisioning and optimization capabilities in IBM Spectrum® Control, some configuration is required.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

To provision storage, you are required to specify only the storage capacity and storage quality that is required. After volumes are created, IBM Spectrum Control can analyze and optimize volume performance. To take advantage of these capabilities, create service classes and, optionally, capacity pools.

A tutorial of cloud configuration concepts is available in the GUI. In the menu bar, go to Advanced Analytics >> Cloud Configuration, and click Learn the Concepts.

- **[Service classes](#)**
A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.
- **[Capacity pools](#)**
Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available capacity for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

Service classes

A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and

whether storage resources must be able to encrypt or thin provision volumes.

Service classes simplify provisioning requests by representing a level or type of storage quality. When you are requesting storage, only the required capacity and service class must be specified. Before you can provision storage, you must create service classes that describe the capabilities and characteristics of the storage you want to be able to provision. Service classes can later be modified or deleted as the needs of your installation change.

Block-storage service classes

A block-storage service class describes attributes and capabilities of block storage resources. When you provision volumes, you specify the requirements by using a block-storage service class. Based on the requirements of the service class, IBM Spectrum® Control identifies a storage pool for the volume.

File-storage service classes

A file-storage service class describes attributes of file storage resources. The file storage resources include Network Attached Storage (NAS) filers, Network Shared Disks (NSDs), and file systems. When you provision shares, you specify the requirements by using a file-storage service class. Based on the requirements of the service class, IBM Spectrum Control identifies a file system or NSD for the share. Space on an NSD is allocated if a new file system is required or an existing file system requires more space.

Service class properties

A service class represents storage resources with common traits. These common traits are the service class properties, and collectively describe a particular quality of service. Before you create service classes, consider the levels of storage quality that you want to have available for provisioning requests. For block storage, for example, you can organize service classes by storage tier or RAID level. Some service class properties are used during provisioning to identify the best location for storage placement. Other service class properties determine how the storage is configured, and reflect the performance, reliability, or security considerations for the level of service.

In addition to the standard properties of a service class, you can create your own properties by using custom tags. Custom tags can represent any common trait that storage systems must have. For more information about creating your own properties by using custom tags, see [Custom tags](#).

Block-storage service classes can specify the following properties:

- Storage tier or range of storage tiers
- RAID level
- Whether volume virtualization is required or not allowed. This property is available only when at least one IBM® storage virtualizer is managed by IBM Spectrum Control. The IBM storage virtualizers include SAN Volume Controller, Storwize® V7000, and Storwize V7000 Unified. If volume virtualization is required, the VDisk mirroring property specifies whether volumes are mirrored to a second storage pool, which uses different backend storage, on the same IBM storage virtualizer.

- Whether volumes must be thin provisioned or must not be thin provisioned. When a thin-provisioned volume is created in a DS8000®, an extent space efficient (ESE) volume is created.

If thin provisioning is required, you can specify thin-provisioning configuration properties for the IBM storage virtualizers and IBM XIV® Storage System. The following properties might be available:

- If the volume is created on an XIV, the locking behavior configuration property affects the volume that is created.
- If the volume is created on an IBM storage virtualizer, the following configuration properties affect the volume that is created:
 - Used capacity
 - Auto-expand
 - Warning level
 - Grain size

- Whether volumes are compressed. This property is applied only when an IBM storage virtualizer has at least one I/O group that contains a compressed volume. If there are no existing compressed volumes in an I/O group, IBM Spectrum Control will not create compressed volumes.

This property is available only when the thin provisioning property is enabled.

- A multipathing policy. If a server has a Storage Resource agent (SRA) and uses an IBM System Storage® Multipath Subsystem Device Driver (SDD), the server multipathing policy is used to configure the driver. If a multipathing policy is set, and fabrics are managed by IBM Spectrum Control, the following properties affect how the storage system is connected to the host:
 - Whether the storage system and the client host are connected through fully redundant paths
 - Number of paths between the storage system and the client host
- If at least one DS8000 Encryption Group is discovered by IBM Spectrum Control, the following properties are available and affect the storage placement and configuration of volumes:
 - Whether volumes must be encrypted
 - If volumes must be encrypted, the DS8000 encryption group to use

File-storage service classes have the following properties:

- Whether dedicated storage is required, or whether shared storage is allowed. In other words, whether the file system from which the NAS share is provisioned can contain other NAS shares.
- Fileset type. The fileset type is either an independent fileset with its own allocated nodes, or a dependent fileset allocated on the file system.
- Access path host name replacement.

Predefined service classes

To illustrate service classes, a set of predefined service classes are provided. From the Service Classes page in the web-based GUI, you can view any of these service classes for more information. Select a service class and then select Actions, > View / Modify. To display help information about a service class property, complete the following steps:

1. Move the mouse pointer over the field or control. A question mark icon is displayed next to the field.

2. Move the mouse pointer over the question mark icon to display an explanation of the property.

The following sample block-storage service classes are provided for illustration.

Gold

The Gold service class represents the highest-performing storage resources for mission-critical applications.

Silver

The Silver service class represents high-performing storage resources for applications in production.

Bronze

The Bronze service class represents standard storage resources for non-mission-critical applications.

The following sample file-storage service classes are provided for illustration.

NormalIsolation

The NormalIsolation service class describes normal isolation file storage by specifying that shared storage is allowed. The file system from which the NAS share is provisioned can contain other NAS shares.

EnhancedIsolation

The EnhancedIsolation service class describes enhanced isolation file storage by specifying that dedicated storage is required. The file system from which the NAS share is provisioned cannot contain other NAS shares.

Custom tags

During provisioning, IBM Spectrum Control determines candidates for storage placement by comparing the requirements of a service class with the known capabilities, configuration, and performance of the available storage systems. You can think of the storage placement determination as a filtering process. IBM Spectrum Control filters the set of available storage resources against properties of the service class until only the storage resources that can satisfy all of the requirements remain.

In addition to the standard properties of a service class, you can create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.

When you specify tags on a block-storage service class, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

When you specify tags on a file-storage service class, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

During provisioning, IBM Spectrum Control filters candidates for provisioning against the custom tags.

Before you create service classes, consider whether there are any special requirements for provisioning that are not addressed by the standard properties of the service class. If so, tag the appropriate storage resources to satisfy custom requirements of a service class. When you are defining the service class, specify the custom tags that are required.

Adding users to a service class

By default, only administrators can provision storage. However, an administrator can add users to a service class to grant them permission to provision by using the service class. In order for an administrator to grant permission to provision to a user, the user must be assigned to the Monitor or External Application role.

When you are granting users permission to provision by using a particular service class, you can also specify whether execution of provisioning plans created by these users requires administrator approval. The requirement for administrator approval applies to all non-administrative users who are added to the service class, and also applies to external applications. Approval cannot be required for some non-administrative users and not for others.

Before you create service classes, consider whether you want to allow users to request their own storage. Consider whether you want user provisioning requests to require administrator approval.

Associating capacity pools with a service class

Before you create service classes, consider whether certain storage requests must always be satisfied from a particular set of resources. If so, you can associate one or more capacity pools with a service class. For example, suppose that a set of storage resources are allocated to a particular department of your business. All storage requests for that department must be satisfied by those storage resources. In this case, you can add all the storage resources to a capacity pool and associate that capacity pool with a service class. You can also add the members of the department to the service class. The department members can then request their own storage as needed. However, the requests can be satisfied only by the storage resources that are allocated to the department.

When you are configuring capacity pools and service classes, keep in mind that a storage resource can be a member of only one capacity pool. Adding storage resources to a capacity pool is not a viable approach if the resources are already members of another capacity pool. Consider specifying custom tags to achieve the same result. In the preceding example, suppose that some or all of the storage resources that are allocated to the department are already members of capacity pools. In this case, you can instead tag the storage systems and the service class with matching tags to achieve the same result. With either approach, only the storage systems that are allocated to the department are candidates for provisioning when storage is requested by using the service class.

Capacity pools

Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. Configure capacity pools to track the used and available capacity for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

The following types of storage resources can be grouped into capacity pools:

- Storage systems
- Storage pools
- File systems of file storage systems

You can organize storage resources into capacity pools in any way that serves your business needs. For example, you can separate the storage resources that are allocated for separate divisions of your business into separate capacity pools. You can then track the storage use for each division separately, and restrict provisioning requests to the appropriate set of storage resources.

Tracking used and available capacity in a capacity pool

You can track the used and available capacity for any set of storage resources by adding the storage resources to a capacity pool. From the Capacity Pools page of the GUI, you can view the total capacity of the resources, and monitor the following capacity measurements:

- Used file space
- Available file space
- Used block space
- Available block space

Restricting a provisioning request to a capacity pool

You can use capacity pools to define a set of storage resources from which provisioning requests must be satisfied. When you are provisioning storage, you can specify a capacity pool. If you do, the provisioning request is restricted to resources in the capacity pool. Only those resources are candidates for provisioning.

You can also associate a service class with capacity pools. If a service class is associated with capacity pools, provisioning requests for the service class must specify, and be constrained to, one of the associated capacity pools. Because you can also grant non-administrative users permission to provision storage by using the service class, associating capacity pools with the service class restricts the users to a specific set of resources.

Provisioning storage overview

IBM Spectrum® Control guides you through the steps for provisioning storage volumes or network-attached storage (NAS) shares. You can provision storage volumes or NAS file shares to one or more servers, one or more hypervisors, or one cluster.

Discontinued support: Provisioning is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your provisioning needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

Storage requirements

When you set up provisioning, you must specify your storage requirements by defining a service class. Certain properties of the service class describe capabilities that storage resources must have so they can provide the service class. For example, when you provision volumes, a block-storage service class specifies properties such as a storage tier, a RAID level, and whether the pool must be able to encrypt or thin provision volumes.

A service class typically represents a quality or service. For example, IBM Spectrum Control provides predefined block-storage service classes that are named Gold, Silver, and Bronze. Gold represents the highest-performing storage resources, and Silver and Bronze represent lesser levels of storage quality.

The different levels of storage quality that are provided by each service class are defined by the service class properties. For example, each service class specifies a different required storage tier that depends on the level of service that is required. Because the Gold service class is for mission critical applications, it sets off thin provisioning.

When IBM Spectrum Control guides you through the steps for provisioning, you concern yourself only with which service class is needed for the new volumes or shares, and how much capacity you require. However, if you are unfamiliar with the service classes, you can open a separate window to view service class information. If you have Administrator privileges, you can modify or create service classes.

Storage constraints

IBM Spectrum Control identifies the storage resources that can provide the capacity and the service class from a set of storage resources. This set might be all the storage resources that are known to IBM Spectrum Control, or it might be constrained to a subset of those resources. In particular, your site can create capacity pools. *Capacity pools* are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. For example, a capacity pool might contain the storage resources that are allocated to a particular department or division of your business.

A service class might restrict storage placement to one or more capacity pools. If the service class you are provisioning from restricts placement, you must select one of the allowed capacity pools. If the service class does not restrict storage placement, you can specify that candidates for provisioning can be selected from all storage resources that are known to IBM Spectrum Control, all storage resources in any capacity pool, or storage resources in a particular capacity pool. If you are unfamiliar with the capacity pools that you can select when you set up provisioning, you can open a separate window to view capacity pool information. If you have Administrator privileges, you can modify or create capacity pools.

If the fabrics are managed by IBM Spectrum Control, the set of storage resources might be further constrained by your zoning policy. If the zoning policy is not configured to automatic zoning, only storage systems with connectivity to the server or hypervisor are candidates for storage placement.

Provisioning volumes

Restriction: To provision volumes, you must have the IBM Spectrum Control Advanced Edition license.

To provision volumes, you must select one or more servers or hypervisors. You can request one or more volumes. For each volume, you specify a volume name, the capacity that is required, a service class, and, optionally, a capacity pool. For each volume, IBM Spectrum Control identifies the storage pools that can provide the capacity and the service class. From the set of pools that can provide the capacity and service class, IBM Spectrum Control identifies the best location for the storage. The best location for the storage is based on the available capacity in the pool and performance data. Preference is first given to storage pools and systems that already contain volumes for the selected server or hypervisor. Preference is then given to systems that have available performance data.

On storage systems, volumes are allocated in increments of a set unit size that depends on the storage system type. Depending on the storage system that was identified as the best location for the storage, the volumes that are allocated might be larger than the capacity you requested for the following reasons:

- On SAN Volume Controller, Storwize® V7000, and Storwize V7000 Unified, volumes are allocated in increments of an extent size. The extent size is set when the MDisk group is created. The capacity that you requested is rounded up, if necessary, to the full extent size.
- On a DS8000®, volumes are allocated in increments of a 1 GiB fixed extent size. The capacity that you requested is rounded up, if necessary, to the full extent size.
- On an XIV®, volumes are allocated in fixed increments of 16 GiB. The capacity that you requested is rounded up, if necessary, to complete the 16 GiB increment.

Provisioning shares

To provision a share, you must first select one or more servers or hypervisors. You specify the capacity that is required for the share, a service class, and information about how to export the share. IBM Spectrum Control identifies the file systems and Network Shared Disks (NSDs) that can provide the capacity and the service class. From the set of file systems and NSDs that can provide the capacity and service class, IBM Spectrum Control identifies the best location for the storage. The best location for the storage is based on the available capacity on the file system or NSD.

The provisioning task

When you complete the steps for setting up provisioning, IBM Spectrum Control creates a provisioning task for the share or volume. If you request multiple volumes, a separate task is created for each unique service class and capacity pool combination you specify in your volumes request. If a requested volume is not constrained to a capacity pool, the set of all available storage resources is considered the capacity pool. If you have Administrator privileges, you can save, run, or schedule provisioning tasks.

In the service class, an administrator can grant users permission to provision by using the service class. In the service class, an administrator can also specify whether scheduling or running provisioning tasks that are created by using the service class requires administrator approval. If you are a user with permission to provision by using the service class, and administrator approval is not required, you can save, run, or schedule the task. You can save the task, if administrator approval is required. Also, if you save the task, an administrator can later schedule or run it.

Although IBM Spectrum Control identifies the best location for storage when it creates the provisioning task, the implementation of the task might fail. For example, the implementation of the task might fail because of changes in the environment that occur after the task is created and before it is run. The implementation of the task might also fail because of storage system restrictions that are not considered when the task is created.

Viewing task details and logs

When a provisioning task is running, you can view its status in a details page. If you run the provisioning task immediately after you complete the steps for setting up provisioning, the details page is already displayed. You can also display details of a provisioning task from the Tasks page. From the details page, you can open logs for the provisioning task. The logs show the steps that are taken by the task during processing and include detailed information about any warnings or errors.

Tracking

When you set up provisioning, you can specify a ticket identifier for tracking purposes. The ticket identifier is associated with the provisioning task, and with any volume or share that is created by the provisioning task.

The ticket identifier can be viewed by showing the Ticket column in the Volumes page or the Shares page. The ticket identifier can also be viewed in the properties notebook for a volume or share.

Tip: If you request multiple volumes, a separate provisioning task is created for each unique service class and capacity pool combination you specify in your volumes request. By specifying a ticket identifier when you request the volumes, you can easily track the set of provisioning tasks and volumes that are created by the single request.

Storage configuration

Whether you are provisioning volumes or shares, IBM Spectrum Control configures the storage and resources as specified by certain properties of the service class, and, in the case of volumes, according to your zoning policy.

When you provision volumes, IBM Spectrum Control can configure volumes and resources according to properties of the service class:

- If the service class requires thin provisioning, the volume is created in a pool that can thin provision volumes. During provisioning, other properties of the service class are used to configure the thin provisioned volume for the particular type of storage system that contains the pool. Because you configure thin provisioned volumes differently for the different storage system types, the service class has a unique set of thin provisioning properties for each type.
- If the service class specifies a multipathing policy for servers, the IBM® System Storage® Multipath Subsystem Device Driver (SDD) on the server is configured to use that policy. For the SDD to be configured, a Storage Resource agent (SRA) must be running on the server.

When you are provisioning volumes, if automatic zoning is enabled, new zones might be created to connect a server to the storage system. Existing zones are used if the server already has connectivity to the storage system. Otherwise, one or more new zones are created between a host initiator port and a controller, node, or module port.

When you provision shares, IBM Spectrum Control configures at least one file access protocol for the file share.

Storage reclamation

You can use volume reclamation recommendations to reclaim unused storage capacity in your environment and use your storage more efficiently.

The block storage systems in your data center that you add for monitoring are regularly analyzed. You can view a list of the volumes that are not mapped to servers or that have no recorded I/O activity. By reclaiming the volumes that are not being used, you can recycle existing storage space instead of purchasing new storage media.

Supported storage systems: Go to see <https://www.ibm.com/support/pages/node/388393>, and then click the release number in the Storage column.

Related tasks

- [Reclaiming storage](#)

Alerts and alert policies

Specify conditions that trigger alerts and the actions to take when those alerts are triggered, such as notify an email address. Use alert policies to define those alert conditions and notification settings for a group of resources.

Alerting functions examine the attributes, capacity, and performance of resources. If the conditions that are defined for alerts are met, the actions that are specified for the alert are taken. Typically, the actions include sending a notification. For example, if the status of a SAN Volume Controller storage system changes to Error, an alert is displayed in the Alerts page in the GUI, and an email might be sent to a storage administrator.

You can manage alerts in your storage environment in the following ways:

- Use alert policies to manage the alert definitions and notification settings that apply to different sets of resources. For example, you can use one alert policy for the storage systems in your test environment, and another for the storage systems in your production environment. Alert policies manage one type of resource only. For example, if you have SAN Volume Controller and FlashSystem 900 storage systems in your storage environment, you cannot have both types of resource in one alert policy.

A resource can be managed by only one alert policy. When you add a resource to be monitored by IBM Spectrum® Control, it is added to a default alert policy automatically.

If a resource is managed by a policy, the resource cannot have alert definitions and notification settings that are independent of the policy. The alert definitions and notification settings that apply to the resource come from the policy.

Default policies with alerts already configured are available. You can create copies of the default policies and assign resources to the new policies. Your alerts are configured with the default settings.

- Define alert conditions and notification settings for individual resources. It is not a requirement for resources to be managed by an alert policy. A resource can have its own alert definitions and notification settings, independent of an alert policy.
- Define alerts and notification settings for applications and general groups. Use applications or general groups to manage alerts for groups of resource components such as volumes or pools. For example, you might want to define alerts on the response time for volumes in an application, depending on the response time requirements of the application. In this case, it is not useful to configure volume response time thresholds for the entire storage system because the storage system might serve many different applications with different needs.

Triggering conditions for alerts

The conditions that trigger alert notifications depend on the type of resource that you are monitoring. In general, the following types of conditions can trigger alerts:

- An attribute or configuration of a resource changed
- The capacity of a resource fell outside a specified range
- The performance of a resource fell outside a specified range
- The storage infrastructure was changed, such as a new or removed resource
- Data is not being collected for a resource

For example, you can use performance thresholds to be notified when the total I/O rate for storage systems falls outside a specified range. This information can help you identify areas in your storage infrastructure that are over used or under used. IBM Spectrum Control provides many metrics for measuring performance and determining violations of the thresholds that you specify.

Alert notifications and triggered actions

When an event occurs and triggers an alert, the alert is written to a log. You can also select one or more other ways to be notified of the event. These alert notifications include SNMP traps, IBM® Tivoli® Netcool®/OMNIBus events, login notifications, entries in Windows event log or UNIX syslog, and emails. Additionally, if a Storage Resource agent is deployed on a monitored server, you can run a script or start an IBM Spectrum Protect job in response to the alert.

Acknowledging alerts: Some alerts are triggered by conditions that commonly occur and can be ignored. In such cases, you acknowledge these alerts to indicate that they were reviewed and do not require immediate resolution. By acknowledging alerts, you can more quickly identify other alerts that must be reviewed and resolved.

Event processing

Alerts are generated when particular conditions are detected during data collection and event processing. For some storage systems such as IBM Spectrum Accelerate and the XIV®, events are polled every minute from the resource.

For other resources, events are subscription-based, where the resource itself or a data source such as an SMI-S provider (also called CIM agent or CIMOM) sends the events to IBM Spectrum Control when conditions change on the resource. Examples of storage systems that use subscription-based event processing include SAN Volume Controller, Storwize® V7000, Storwize V7000 Unified, and FlashSystem V9000. For these storage systems, a probe is automatically run when many events are received from the storage system in a short time period. To avoid performance bottlenecks, probes are run only every 20 minutes.

Prerequisites for using alerts

The following conditions must be met to successfully use alerts:

- Data collection schedules are configured and scheduled to run regularly. For example, to detect violations of performance thresholds, you must run performance monitors to collect performance data about resources. Running performance monitors regularly also helps to establish a history of performance for trending analysis.
- If you want to be notified about an alert in some way other than an entry in the log file, such as using SNMP traps, IBM Tivoli Netcool/OMNIBus events, or email, you must configure those alert destinations before you use the alert.
- If an alert is triggered based on an SNMP trap from the monitored resource, you must properly configure the SNMP server of the monitor resource to enable IBM Spectrum Control to listen to SNMP traps. The default port number is 162, and the default community is public.

Related concepts

- [Collecting data](#)
- [Alerting](#)

Related reference

- [Performance metrics](#)
- [Triggering conditions for alerts](#)
- [Alert notifications and actions](#)

Reporting

Use reporting functions to view overview and detailed information about your storage.

You can view detailed capacity information for storage systems, servers, hypervisors and their internal resources, such as volumes, pools, disks, and virtual machines. For example, you can view the total amount of storage capacity that is committed to a volume, the capacity that is used by a volume, and the capacity that is allocated to a volume and is not yet used.

In rollup mode, secondary rollup servers collect information from managed resources and roll that data up to primary rollup servers. This gives you a network-wide perspective of storage usage in your environment when you have multiple IBM Spectrum® Control servers deployed.

You can use the IBM® Cognos® Analytics reporting tool to view predefined reports and create custom reports about the resources managed by IBM Spectrum Control. Use the IBM Cognos Analytics reporting tool to view predefined reports about the capacity and performance of your resources. Charts are automatically generated for most of the predefined reports. Depending on the type of resource, the charts show statistics for space usage, workload activity, bandwidth percentage, and other statistics. You can schedule reports and specify to create the report output in HTML, PDF, and other formats. You can also configure reports to save the report output to your local file system, and to send reports as email attachments.

You must collect information about your environment before you can use reports to view details about the storage resources in it. You can use IBM Spectrum Control monitoring jobs, such as probes and performance monitors, to gather comprehensive information and statistics about your storage resources.

Related concepts

- [Reporting with Cognos Analytics](#)

Related reference

- [Rollup servers](#)

Rollup servers

Rollup reporting combines capacity and status information from multiple instances of IBM Spectrum® Control for monitoring and reporting. In rollup mode, secondary rollup servers collect information from managed resources and roll that data up to primary rollup servers.

Rollup capability requires two types of servers: a primary rollup server running IBM Spectrum Control, which collects capacity and status information from one or more secondary rollup servers. Secondary rollup servers communicate with the primary rollup server during server probes.

The primary rollup server collects capacity and status information about storage resources that are managed by the secondary rollup servers. This primary/secondary rollup server configuration is called *rollup mode*. When you enable rollup mode on the primary rollup server, you can see capacity and status information in the rollup resource pages for the storage systems, servers, hypervisors, switches, and fabrics that are monitored by your rollup servers.

You can also configure a primary rollup server as a secondary rollup server to other primary servers. This allows you to create a comprehensive rollup reporting structure across your enterprise.

Note: If you configure a primary rollup server as a secondary rollup server to another primary server, the data collected by that primary server is not rolled up to the higher primary server. Primary rollup servers only collect and display data from their immediate secondary rollup servers. For example, if you configure primary rollup server 2 as a secondary to primary rollup server 1, the data that is rolled up to primary rollup server 2 from its secondary servers is not collected by primary rollup server 1. Primary rollup server 1 only collects capacity and status information about primary rollup server 2, not the secondary servers monitored by primary rollup server 2.

Note: After you upgrade a primary rollup server you must run probes of the secondary rollup servers that were added to the primary server prior to the upgrade. Use the Start Probe action to run probes of the secondary rollup servers. Wait for the first probe (manual or scheduled) to complete, before you work with a new function.

Managing secondary rollup servers

After you add a secondary rollup server on the Settings > Rollup Server Connections page, you can view information about the server and perform actions on it. You must have IBM Spectrum Control Administrator role authorization to add, remove, and manage secondary rollup servers on this page.

Information about secondary rollup servers

IBM Spectrum Control provides the following information about the secondary rollup servers on the Rollup Server Connections page:

Managed resources

The number of block storage systems, servers, switches, fabrics, and hypervisors that the secondary rollup server is managing. You can view information about these resources on the rollup resource list pages for the resources, such as the Switches (Rollup) page. NPV switches are not included in this number on either the primary rollup or secondary rollup Switches (Rollup) resource list pages.

Data collection

The status of the last probe of the secondary rollup server, and the time of the next probe of the server. Hourly probes are automatically scheduled when a secondary rollup server is added to a primary rollup server.

Actions

You can perform the following actions on the secondary rollup server. Unless otherwise noted, you must have Administrator role authority to perform these actions.

Open logs

Opens the probe logs window for the secondary rollup server in a separate browser window. Any IBM Spectrum Control user role can perform this action.

Start probe

Starts a probe of the secondary rollup server. If a probe is already running for the server, you must wait until the current probe completes before starting a new probe.

Remove server

Removes the secondary rollup server from the Rollup Server Connections page and the Servers (Rollup) list. Also removes any associated resources that are managed by the secondary rollup server. For example, if a secondary rollup server is the managing instance for a switch and the secondary server is removed from the primary rollup server, the switch will no longer display on the Switches (Rollup) resource list page.


Note: Removing a secondary rollup server on the Rollup Server Connections page does not remove that server from the non-rollup Servers page. If the server has been added by using the Add Server dialog on the non-rollup Servers page, IBM Spectrum Control will continue to manage that server unless you remove it from the Servers page.

Modify connection

Allows you to change the host name, user name, password, and port of the secondary rollup server.

Rollup mode

Rollup mode is a special view of your enterprise and is enabled by configuring your IBM Spectrum Control servers to communicate and gather data through server probes. Rollup mode requires a primary rollup server configured to collect capacity and status information from one or more secondary rollup servers. Once you have configured a


primary rollup server with at least one secondary rollup server, you can go to the rollup icon  in the menu bar to enter rollup mode. The rollup icon does not display until you have configured at least one secondary rollup server for the primary rollup (local) server. You might have to refresh your browser window to see the rollup icon after you add the first secondary rollup server for the primary rollup server you are configuring.

Entering rollup mode

To view rollup capacity and status information from your IBM Spectrum Control servers, you must enter rollup mode.

1. On the Rollup Server Connections page, add at least one secondary rollup server. Secondary rollup servers must have IBM Spectrum Control installed and running.
2. Move the mouse pointer over the rollup icon in the menu bar and select Enter rollup mode.

You will see the following changes in the menu bar:

- The color of the rollup icon changes: .
- The lower edge of the menu bar changes to orange.
- The menu bar options change to Storage, Servers, Network, Rollup Server Connections.

Using rollup mode

Rollup mode provides lists of resources that are managed by secondary rollup servers. For example, you can see a list of all the network switches that are managed by a primary rollup server and its associated secondary rollup servers. These resource list pages are similar to the pages you work with when not in rollup mode. However, in some cases rollup mode provides different data columns and different actions that you can take on the rollup storage devices.

View rollup information

1. In the rollup menu bar, select the type of storage device you want to view information about. For example, Servers.
2. Select one of the available menu options. The Servers menu bar item has two options: Servers and Hypervisors.
3. On the resource list page, you can see the data about that resource type that has been collected from the primary and secondary rollup servers.
4. Rollup resource lists function similar to non-rollup pages: you can add and reorder data columns, right-click items in the resource list to perform actions on them, and other familiar functions. You can also use Ctrl-click to select multiple list items.

Actions

You can perform actions on items in the resource lists by right-clicking them or by selecting one or more list items and clicking Actions.

View properties

View the properties of the selected item in the resource list. This action is not available for multiple selections at the same time.

Export

Export the data for one or more list items. You can specify PDF, CSV, or HTML output.

Units of measurement for storage data

IBM Spectrum® Control uses decimal and binary units of measurement to express the size of storage data.

Decimal units such as kilobyte (KB), megabyte (MB), and gigabyte (GB) are commonly used to express the size of data. Binary units of measurement include kibibyte (KiB), mebibyte (MiB), and gibibyte (GiB). [Table 1](#) compares the names, symbols, and values of decimal and binary units.

Table 1. Comparison of binary and decimal units and values

Binary			Decimal		
Name	Symbol	Value (base 2)	Name	Symbol	Value (base 10)
kibibyte	KiB	2 ¹⁰	kilobyte	KB	10 ³
mebibyte	MiB	2 ²⁰	megabyte	MB	10 ⁶
gibibyte	GiB	2 ³⁰	gigabyte	GB	10 ⁹
tebibyte	TiB	2 ⁴⁰	terabyte	TB	10 ¹²

Binary			Decimal		
Name	Symbol	Value (base 2)	Name	Symbol	Value (base 10)
pebibyte	PiB	2 ⁵⁰	petabyte	PB	10 ¹⁵
exbibyte	EiB	2 ⁶⁰	exabyte	EB	10 ¹⁸

Binary units of measurement express the size of data more accurately. When you compare the size of 100 KB to 100 KiB, the difference is relatively small, 2.35%. However, this difference grows as the size of the data values increases. When you compare the size of 100 TB to 100 TiB, the difference is 9.06%.

In general, IBM Spectrum Control uses base 2 values for memory and disk space values, and base 10 values for space on physical hard drives.

[Table 2](#) shows the percentage difference between decimal and binary values across a range of data sizes.

Table 2. Percentage difference between decimal and binary units

Decimal value	Binary equivalent of decimal value	Difference
100 kilobytes (KB)	97.65 kibibytes (KiB)	2.35%
100 megabytes (MB)	95.36 mebibytes (MiB)	4.64%
100 gigabytes (GB)	93.13 gibibytes (GiB)	6.87%
100 terabytes (TB)	90.94 tebibytes (TiB)	9.06%
100 petabytes (PB)	88.81 pebibytes (PiB)	11.19%
100 exabytes (EB)	86.73 exbibytes (EiB)	13.27%

Role-based authorization


Roles determine the functions that are available to users of IBM Spectrum® Control. When a user ID is authenticated to IBM Spectrum Control through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

The following table shows the IBM Spectrum Control roles and their authorization levels:

Table 1. IBM Spectrum Control roles and authorization levels

Roles	Authorization level
Administrator	This role has full access to all monitoring and administrative functions. At least one group must have the Administrator role. Note: When IBM Spectrum Control is first installed, the following operating system groups are assigned the Administrator role: <ul style="list-style-type: none"> Windows: Administrators UNIX and Linux®: root AIX®: system
Monitor	This role has access to the following read-only functions: <ul style="list-style-type: none"> Viewing and exporting information about monitored resources Viewing, acknowledging, and removing alerts Viewing tasks and data collection jobs Opening management GUIs Opening logs Viewing chargeback, consumer, predefined capacity and inventory, and custom reports Exception: Users with the Monitor role can provision storage if they are granted permission in a service class. A service class is a logical entity that describes storage capabilities and characteristics and can be used to specify requirements for storage provisioning. For more information about service classes, see Creating service classes .
External Application	If you assign the External Application role to the user, you must also assign one or more service classes to the user. This role does not enable users to log in to the IBM Spectrum Control GUI.

Tips:

- To determine the role of the user who is logged in, click the user icon  in the upper-right corner of any page in the GUI.
- If a user belongs to multiple groups and the groups have different roles, the role with the highest level of authorization is granted to the user. For example, if a user belongs to a group that is assigned the Administrator role and also belongs to a group that is assigned a Monitor role, the user is granted the authorization of the Administrator role.
- If a user is not a member of a group that is assigned an IBM Spectrum Control role, no access is granted to that user.
- If assigned the Monitor role, a user can only open and view logs from the Data Collection page for the selected resource.

Nested groups are not supported: Adding active directory or any other type of domain user group to a local operating system group is not supported in IBM Spectrum Control. You can configure IBM Spectrum Control to authenticate domain IDs that rely on the operating system to perform the authentication operation against the active directory, but it cannot resolve nested groups.

Alternatively, you can configure LDAP authentication to perform queries against active directory user repositories and assign domain groups directly to roles within IBM Spectrum Control.

Fabrics and zones

Use the IBM Spectrum® Control to learn more about fabrics and zones.

- [Monitoring fabrics and zones in the GUI](#)

You can monitor switches and fabrics in the IBM Spectrum Control GUI. You can also enable automatic zoning so that new zones are created during storage provisioning, if necessary, to connect a server to a storage system.

- [Zones, zone aliases, and zone sets](#)

Zones, zone aliases, and zone sets allow logical grouping of ports and storage devices within a storage area network. This section describes zoning concepts and elements.

- [Switch zoning capabilities](#)

When you select a fabric in which to perform zoning operations, Fabric Manager determines the capabilities of the switches in the SAN and limits the zoning operations based on that information.

Related tasks

- [Provisioning with zone control](#)

Monitoring fabrics and zones in the GUI

You can monitor switches and fabrics in the IBM Spectrum® Control GUI. You can also enable automatic zoning so that new zones are created during storage provisioning, if necessary, to connect a server to a storage system.

You can use the GUI to add switches and fabrics for monitoring. You can add the following types of switches:

- Brocade
- Cisco

After switches and fabrics are probed, you can view detailed status and performance information about the switches and fabrics. You can also test the connection to the switch or fabric.

Information that you can view about a fabric includes the following properties:

- The fabric type
- The number of physical and virtual switches in the fabric
- The number of ports that are on all the switches in the fabric
- The number of switch ports that are online and connected to other ports in a fabric
- The name of the principal switch of the fabric

Information that you can view about a switch includes the following properties:

- The number of ports on a switch
- The number of ports that are connected to a storage resource
- The name of the fabric where a switch is a member
- The name of the vendor or manufacturer for a switch

From the Details page of a fabric, you can open the Zone Sets page to view information about the zone sets on the fabric, and to identify which is the active zone set.

You can enable automatic zoning in the GUI. When automatic zoning is enabled, IBM Spectrum Control creates new zones during provisioning if new zones are needed to connect the storage system with the server.

Zones, zone aliases, and zone sets

Zones, zone aliases, and zone sets allow logical grouping of ports and storage devices within a storage area network. This section describes zoning concepts and elements.

In a storage area network a *zone* is a logical grouping of ports to form a virtual private storage network. Zones that belong to a single SAN can be grouped into a *zone set*, which can be activated or deactivated as a single entity across all switches in the fabric. A zone set can contain one or more zones, and a zone can be a member of more than one zone set. Using zoning, you can automatically or dynamically arrange fabric-connected devices into logical groups across a physical fabric.

Ports and devices in a zone are called *zone members*. A zone can contain one or more zone members. Ports that are members of a zone can communicate with each other, but they are isolated from ports in other zones. Devices, however, can belong to more than one zone. A *zone alias* is a collection of zone members. A zone alias can be added to one or more zones.

Note: In some cases, inactive zones might not have any zone members. Activating an inactive zone set that contains empty zones will fail if a switch does not support empty zones in active zone definitions.

Zone membership can be specified by:

- The N_Port_Name of the N_Port connected to the switch (also known as WWN zoning or port name zoning)
- The N_Port address identifier assigned during fabric login (also known as PortId or FCID zoning)
- The Node_Name associated with the N_Port
- The Domain identification (Domain_ID) and physical port ID of the Switch Port to which the N_Port is attached (also known as domain port zoning)
- An alias name

Zoning supports the use of aliases, which are meaningful names assigned to devices. An alias can also be a group of devices that are managed together to make zoning easier.

There are two types of zoning:

Hardware zoning (port zoning)

In hardware zoning (also called *port zoning*), the members of a zone are the physical ports on a fabric switch.

Software zoning (WWN zoning)

Software zoning uses the Simple Name Server (SNS) that runs inside a fabric switch. It is based on the node WWN or port WWN of the zone members to be included. Software zoning lets you create symbolic names for the zones and zone members.

A *default zone* is a group of devices that are not members of the active zone set. These can communicate with each other but not with members of any other zone. Default zoning is enabled by default. You can use a switch element manager to configure the Default Zone option to enable or disable the default zone independently of the active zone set.

Note:

1. If the default zone is disabled, devices that are not members of the active zone set cannot communicate.
2. If the default zone is disabled and no zone set is active, no devices can communicate.
3. If default zoning is enabled, deactivating the active zone set makes all devices members of the default zone. If default zoning is disabled, all communication stops.
4. If you activate one zone set while another zone set is active, the currently active zone set is deactivated.
5. If your EFC Manager manages multiple fabrics, ensure that you have the correct zone set for the fabric you are currently updating.

Switch zoning capabilities

When you select a fabric in which to perform zoning operations, Fabric Manager determines the capabilities of the switches in the SAN and limits the zoning operations based on that information.

To see a list of supported switches and their capabilities, default values, ranges, and possible effects, go to [IBM Spectrum Control interoperability matrix for switches](#).

- [Zoning for Cisco MDS 9000 switches](#)
Cisco switches support virtual SANs (VSANs), which is the logical partitioning of a fabric into multiple fabrics. The overall network is referred to as the physical infrastructure, and the logical fabrics are the VSANs. IBM Spectrum Control provides basic zone discovery for the Cisco MDS 9000 series switches.
- [Brocade switches in zones](#)
Brocade switches that are managed by using Brocade Network Advisor support some non-standard zones such as quick loop zones, fabric assist zones, and protocol zones. If the switch configurations have these zones already defined, IBM Spectrum Control preserves them and does not modify them in any way.

Zoning for Cisco MDS 9000 switches

Cisco switches support virtual SANs (VSANs), which is the logical partitioning of a fabric into multiple fabrics. The overall network is referred to as the physical infrastructure, and the logical fabrics are the VSANs. IBM Spectrum® Control provides basic zone discovery for the Cisco MDS 9000 series switches.

IBM Spectrum Control supports these zone member types:

- N_Port WWN
- FC ID (the fibre channel ID of an N_port attached to the switch)

IBM Spectrum Control does not allow zone management of zones that contain unsupported members.

Brocade switches in zones

Brocade switches that are managed by using Brocade Network Advisor support some non-standard zones such as quick loop zones, fabric assist zones, and protocol zones. If the switch configurations have these zones already defined, IBM Spectrum® Control preserves them and does not modify them in any way.

You can create, change, and delete non-standard zones by using the Brocade switch management application.

Note: IBM Spectrum Control allows you to add empty zones to an inactive zone definition. Activation of a zone set containing empty zones may still fail if the switch does not support empty zones in active zone definitions.

Agents

The IBM Spectrum® Control uses agents to gather data: Common Information Model (CIM) agents, Storage Resource agents, and SNMP agents.

- [SMI-S providers](#)
SMI-S providers enable communication between storage devices, such as storage systems and fabric switches, and IBM Spectrum Control. Each storage device that you want to manage must have an SMI-S provider either installed or embedded unless it is a device that uses Native API connection.
- [Storage Resource agents](#)
Use Storage Resource agents to collect information about storage resources such as servers, virtual machines, workstations, and HBAs.
- [SNMP agents](#)
An SNMP agent, also called an out-of-band agent, is software that runs on a switch and can be used to monitor and manage systems and devices in your network. IBM Spectrum Control uses SNMP to retrieve information about specific switches, such as Cisco switches, and the fabrics that the switches are members of.

SMI-S providers

SMI-S providers enable communication between storage devices, such as storage systems and fabric switches, and IBM Spectrum® Control. Each storage device that you want to manage must have an SMI-S provider either installed or embedded unless it is a device that uses Native API connection.

SMI-S providers are provided by the vendor of the storage subsystem or fabric switch. Each vendor provides unique agent code for their family of storage devices. This code implements a Common Information Model Object Manager (CIMOM) that conforms to the Storage Management Initiative Specification (SMI-S) of the Storage Networking Industry Association (SNIA).

The SMI-S provider enables communication between the storage device and IBM Spectrum Control. Commands and responses are transmitted between IBM Spectrum Control and the SMI-S provider using an XML transport layer. The SMI-S provider to storage device layer uses a vendor-specific proprietary interface.

The SMI-S provider usually must be installed and configured, so that it can identify the storage devices with which it communicates. Some storage devices, such as fabric switches, contain embedded SMI-S providers and so do not require that SMI-S providers be installed. In these cases, IBM Spectrum Control must be configured to point directly to the storage devices that contain the embedded SMI-S providers.

SMI-S providers can be referred to as CIM agents, CIM proxy agents, and CIMOM agents. SMI-S providers can be embedded in the device or installed on a separate computer.

Note:

- Do not install multiple SMI-S providers on a single computer because of port conflicts.
- Do not install a SMI-S provider on the system where a IBM Spectrum Control server component is installed.

Related tasks

- [Upgrading CIM agents](#)

Related reference

- [Storage Management Initiative Specification](#)
- [SMI-S providers](#)

Storage Resource agents

Use Storage Resource agents to collect information about storage resources such as servers, virtual machines, workstations, and HBAs.

You must deploy Storage Resource agents on resources where you want to gather the following information:

- Asset information
- File and file system attributes
- Network-attached storage (NAS) device information

Tip: IBM Spectrum® Control creates and updates agentless servers automatically after it probes storage systems and hypervisors.

Related tasks

- [Adding servers with Storage Resource agents](#)

Related reference

- [Deployment guidelines and limitations for Storage Resource agents](#)
- [Planning for Storage Resource agents](#)

SNMP agents

An SNMP agent, also called an out-of-band agent, is software that runs on a switch and can be used to monitor and manage systems and devices in your network. IBM Spectrum® Control uses SNMP to retrieve information about specific switches, such as Cisco switches, and the fabrics that the switches are members of.

IBM Spectrum Control uses SNMPv3 (preferred) or SNMPv1 to probe Cisco switches to collect performance data. Some switches are configured to use SNMPv3 by default.

Related tasks

- [SNMP agents](#)
- [Displaying information about an SNMP agent](#)

Related reference

- [Cisco SNMP agent](#)
- [Simple Network Management Protocol](#)
- [Planning for Cisco](#)

IBM Spectrum Control REST API

You can use the Representational State Transfer (REST) API for IBM Spectrum® Control to access information about resources and to generate custom capacity, configuration, and performance reports.

To get the information that you need about your resources, you can use a REST command line utility or you can use a web browser.

Related tasks

- [Using the REST API to generate reports](#)

Tutorials

Use tutorials to learn how to perform specific tasks and become familiar with IBM Spectrum® Control.

The provided tutorials set out situations that can be solved by using IBM Spectrum Control. When using these tutorials you should remember the following:

- The tutorials are intended to be step-by-step instructions that you follow to complete a specific job. Each tutorial contains multiple tasks that must be followed in the order given.
- Specific values are given in the tutorials, such as user name, IP Address, probe name, and so forth. These are for illustration purposes only and you must replace them with values appropriate for your system.
- The tutorials do not provide in-depth conceptual information about the tasks. See the appropriate topics in the IBM Spectrum Control Knowledge documentation for more information about the concepts behind the tasks.
- You will complete a tutorial successfully if you follow the instructions correctly and your system performs as expected.
- **[Optimizing the performance of storage virtualizers](#)**
In this tutorial, you balance pools and analyze tiering to help resolve performance hotspots and manage storage resources efficiently.
- **[Monitoring capacity usage at different levels of a business hierarchy](#)**
The monitoring and management of applications and departments enables you to monitor storage capacity usage, recognize trends, monitor health status, and troubleshoot performance of the storage resources in your business organization.
- **[Using applications and subcomponents to monitor capacity and space usage](#)**
To monitor the performance, capacity, and space usage of the applications in your business organization, create applications and subcomponents. You can also add applications to departments so that storage capacity and usage can be monitored in an overall business hierarchal manner.
- **[Tutorial: Viewing NPIV connections between server ports and switch ports in a fabric](#)**
In this tutorial, a storage administrator views the N-Port ID Virtualization (NPIV) connections between the server ports and the switch port in the fabric. With NPIV, multiple node ports can be logically connected to one switch port.
- **[Tutorial: Exporting and uploading performance data for a SAN Volume Controller system](#)**
In this tutorial, you export performance data for a SAN Volume Controller to a compressed package » and upload it to a support ticket. « You can also send the package to IBM® Support.
- **[Tutorial: Comparing the performance of storage systems](#)**
In this tutorial, you compare the performance of two storage systems.
- **[Tutorial: Reviewing and updating your agentless servers](#)**
In this tutorial, you review the agentless servers that are created automatically by IBM Spectrum Control, and update your agentless servers.
- **[Tutorial: Troubleshooting performance](#)**
Bob is a storage administrator. One morning, he receives a ticket: A critical medical application, Epic Database, has a performance problem. Bob must investigate the problem and identify the cause.
- **[Tutorial: Monitoring IBM Spectrum Scale performance](#)**
You can use IBM Spectrum Control to monitor the performance of the IBM Spectrum Scale clusters in your storage environment.
- **[Tutorial: Viewing the aggregated workload for an application](#)**
Bob is a storage administrator. Bob wants to easily evaluate the performance of the accounting application because the Accounts department is expanding to a new location and they want to replicate their data across both locations. Bob needs to work out the inter-site link capacity requirements.
- **[Tutorial: Identifying the source of slow drain problems caused by depletion of buffer credits](#)**
Use this tutorial to find out how to use IBM Spectrum Control to identify a host that has depleted buffer credits that are causing a slow drain condition.
- **[Tutorial: Identifying the locations of devices](#)**
Identify the locations of devices that are monitored by IBM Spectrum Control. You can also apply custom tags to more easily identify, sort, or group devices based on location or another attribute that you specify.

Optimizing the performance of storage virtualizers

In this tutorial, you balance pools and analyze tiering to help resolve performance hotspots and manage storage resources efficiently.

Before you begin

You can optimize the placement of volumes on the following storage virtualizers:

- SAN Volume Controller
- Storwize® V7000
- Storwize V7000 Unified block storage

Learn more: For information about the components that can be used with IBM Spectrum® Control Version 5.2 (or later), see [IBM Spectrum Control interoperability matrix](#).

To optimize the placement of volumes in storage virtualizer pools, you must complete the following tasks:

- Probe the storage virtualizers.
- Set the tier level of the storage pools that you want to analyze.
- Collect performance data to analyze the pools and volumes.

When you run the analysis to balance pools, you optimize storage performance by redistributing volume workloads across pools on the same tier.

When you run the tiering analysis, you optimize storage performance by specifying thresholds to move volumes to higher and lower tiers.

- [Tutorial: Analyzing and re-tiering volumes in pools on tier 1](#)

In this tutorial task, you analyze the current placement of volumes in pools on tier 1 so that you can move volumes with low workloads from tier 1 to tier 2 or tier 3 pools.

- [Tutorial: Collocating volumes](#)

In this tutorial task, you want to minimize the exposure of servers to multiple back-end storage systems by collocating volumes that are assigned to the same hypervisor or server. You can enforce the collocation of volumes when you enter the criteria for analyzing tiering and balancing pools.

Related tasks

- [Optimizing storage tiering](#)
- [Optimizing storage pools](#)
- [Managing tasks for tiering storage, balancing pools, and transforming storage](#)

Tutorial: Analyzing and re-tiering volumes in pools on tier 1

In this tutorial task, you analyze the current placement of volumes in pools on tier 1 so that you can move volumes with low workloads from tier 1 to tier 2 or tier 3 pools.

About this task

The pools that you want to analyze are on a SAN Volume Controller with two DS8000® storage systems. The pools on the storage virtualizer are assigned to tier 1, tier 2, and tier 3.

Choosing the performance data that is used in the analysis



1 Analysis period and data collection days

The analysis period is set to 10 days and Saturdays and Sundays are excluded from the analysis.

Mo	Tu	We	Th	Fr	Sa	Su
Mo	Tu	We	Th	Fr	Sa	Su



The performance data that is collected is used in the analysis.




The performance data that is collected isn't used in the analysis.



2 Data collection period

09:00 is selected as the start time and 18:00 is selected as the end time.



3 Analysis schedule

You create the analysis on Monday and schedule the analysis to run every 14 days.

Mo	Tu	We	Th	Fr	Sa	Su
Mo	Tu	We	Th	Fr	Sa	Su
Mo	Tu	We	Th	Fr	Sa	Su



The day that you created the analysis is the day that the analysis is run.

For example, you want to analyze the data that is collected over 10 days and when the volumes are most active. So, you set the period for the analysis to 10 days and exclude the data that is collected at the weekend. You also exclude the data that is collected before 9 AM and after 6 PM because you know that the volumes are most

Tiering and volume placement criteria



1 Tiering criteria

The source volumes that are selected are in pools on Tier 1, and the target pools that are selected are on Tier 1, Tier 2, and Tier 3. The tiering thresholds for the volumes are >1000 I/O per second for Tier 1 and >500 I/O per second for Tier 2.

Tiering Criteria	Tier 1: Volume Analysis	Tiering recommendations
Tier 1 I/O Rate: >1000	 I/O Rate = 900	 Place volume in pool on Tier 2
Tier 2 I/O Rate: <1000 and >500	 I/O Rate = 400	 Place volume in pool on Tier 3
Tier 3 I/O Rate: <500		



2 Volume placement criteria

To ensure that the pools that are selected as target pools can handle the additional workload, you specify the maximum I/O rates for the target pools on each tier.

Target Pools	Below Threshold	Add Volume
 Target Pool 1	 Yes	 Yes
 Target Pool 2	 No	 No

You want volumes with the highest workloads in tier 1 pools and you want volumes with low workloads in tier 2 or tier 3 pools. To tier your storage, you enter two I/O rate thresholds to generate recommendations to re-tier the volumes.
Tip: You can tier volumes by I/O rate or I/O density.
To ensure that the pools that are selected as destination pools for the volumes can handle the additional workload, you specify the maximum I/O rates for pools on each tier of storage. Only the pools with I/O rates below the maximum I/O rate can be selected as destination pools for the volumes.

Procedure

1. From the Storage menu, click Storage Systems.
2. Right-click the storage virtualizer and click View Details.
3. In the navigation pane, click Pools.
To organize the pools by tier level, click the Tier column heading.
4. Select the pools on tier 1, right-click and then click Analyze Tiering.
5. Select all of the pools on tier 2 and tier 3 as target pools, and then click Next.
6. Select the period that you want to use to analyze the performance data.
7. Set the threshold for moving volumes from tier 1 to tier 2.
The volumes with I/O rates above the threshold that you set remain in the tier 1 pools. The volumes with I/O rates below the threshold that you set are moved down to tier 2 pools.
8. Set the threshold for moving volumes from tier 2 to tier 3.
The volumes with I/O rates above the threshold that you set remain in the tier 2 pools. The volumes with I/O rates below the threshold that you set are moved down to tier 3 pools.
9. Set the maximum I/O rate for pools on each tier.
10. Click Analyze.

Results

The Analyze Tiering task is shown.

What to do next

Click Execute to implement the recommendations. To create a schedule, click [Schedule > Analysis](#). Alternatively, you can modify the criteria that you set for analyzing the volumes on the Tasks page and run the analysis again.

Related tasks

- [Modifying the criteria for analyzing tiering](#)

Tutorial: Collocating volumes

In this tutorial task, you want to minimize the exposure of servers to multiple back-end storage systems by collocating volumes that are assigned to the same hypervisor or server. You can enforce the collocation of volumes when you enter the criteria for analyzing tiering and balancing pools.

About this task

To ensure that volumes in the same storage pool that are assigned to the same server or hypervisor are kept together, you want to enforce the collocation of volumes. By enforcing the collocation of volumes, you prevent the placement of related volumes in destination pools that might be on multiple back-end storage systems. Multiple host connections to the same hypervisor or server: If the volumes in the source pool that are assigned to the same hypervisor or server are assigned to different host connections, the collocation of the volumes is affected. In such cases, if volumes require optimization, the volumes that are assigned to the same host connection are kept together. To view information about the host connection for the volume, right-click the volume, select View Details, and then click the Host Connections tab.

Procedure

1. From the Storage menu, click Pools.
2. Right-click the pools that you want to analyze, and then click Analyze Tiering.
3. Select the target storage pools.
4. On the Optimize the Placement of Volumes page, ensure that Collocate volumes is set to Yes.
5. Click Analyze.

Results

If a volume requires re-tiering, a recommendation is generated to move all of the volumes that are assigned to the same server to the same destination pool. If none of the target pools have sufficient space to accommodate all of the volumes, then recommendations to move the volumes are not generated.

Related tasks

- [Optimizing storage pools](#)
- [Optimizing storage tiering](#)
- [Modifying the criteria for analyzing tiering](#)

Monitoring capacity usage at different levels of a business hierarchy

The monitoring and management of applications and departments enables you to monitor storage capacity usage, recognize trends, monitor health status, and troubleshoot performance of the storage resources in your business organization.

Before you begin

The department model comprises these main elements:

- The department
- The subdepartments that the department contains
- The applications, or the application subcomponents, or both that the department uses

In the following tutorial, you want to create a department model that monitors the capacity and space usage of the Books Sales department, and these subdepartments:


- Wholesale
- Retail
- Online

To monitor the capacity and space usage of the departments, you add the applications that the department and its subdepartment use to each of the subdepartments.

The department model that you create determines how you can view capacity usage and space information. To view the total capacity and space usage of a department regardless of the subdepartments that the department contains, you create a department and add the applications and subcomponents to the department.

If, however, you want to see the capacity and space usage of the subdepartment and its subdepartments, you create a department hierarchy by adding subdepartments to departments. You then add the applications and application subcomponents that each subdepartment uses to the subdepartments. You can then view capacity and space usage for the department and for the individual subdepartments.

You can create more complex models of departments by adding subdepartments to subdepartments. In the following illustration, the department model is extended to include two more layers of subdepartments under the Wholesale subdepartment.



To monitor the capacity and space usage of the departments in your business organization, complete the following tasks:

- Create the Book Sales department.
- Create the Wholesale, Retail, and Online departments to represent the subdepartments.
- Add the Wholesale, Retail, and Online as subdepartments to the Book Sales department.
- Create the Book Sales DB application and the Wholesale Transactions, Retail Transactions and Online Transactions applications to represent the subcomponents that are used by the departments.

Different applications require different storage capabilities, and different levels of performance and uptime. The storage requirements of the departments are constantly growing and they need to modify their behaviors in a way that justifies the cost of their actions concerning storage utilization. The data modeling of storage resources using the application and department concept enables you to plan and implement a chargeback system if necessary.

- [Tutorial: Comparing storage usage in each department](#)
In this tutorial, you compare the storage usage in each department within your organization to spot a potential storage usage issue.

Related concepts

- [Departments](#)

Related tasks

- [Creating departments](#)
- [Adding and creating subdepartments](#)
- [Viewing information about departments](#)
- [Creating applications](#)
- [Adding subcomponents](#)

Tutorial: Comparing storage usage in each department

In this tutorial, you compare the storage usage in each department within your organization to spot a potential storage usage issue.

About this task

To monitor the capacity and space usage of a department, you create the department and subdepartments, and you add the application subcomponents to the subdepartments.

You want to monitor the capacity and space usage of departments in an organization that sells books.

The sales department of the organization, Book Sales, has these subdepartments:

- Wholesale
- Retail
- Online

Although you can add applications and application subcomponents when you create departments, it is easier and quicker to create the applications and application subcomponents beforehand.

You want to know the total amount of storage space that the Book Sales department uses, and you want to know the amount of space that each subdepartment uses. To know how much space the Book Sales and its subdepartments use, you must associate the department and its subdepartments with the applications that they use. For example, the Book Sales department and its subdepartments use a database application named Book Sales DB. The Book Sales DB application contains these subcomponents:

- Wholesale Transactions
- Retail Transactions
- Online Transactions

Procedure

To create the department and subdepartments, you complete these tasks:

1. You create the Wholesale department and add the applications and application subcomponents that the subdepartment uses.
2. You create the Retail department and add the applications and application subcomponents that the subdepartment uses.
3. You create the Online department and add the applications and application subcomponents that the subdepartment uses.
4. You create the Book Sales department.
You do not have to add applications or subcomponents to the Book Sales department because the Book Sales department inherits the applications and application subcomponents that are added to the subdepartments.
5. You add the Wholesale, Retail, and Online departments as subdepartments to the Book Sales department.

Results

When you complete adding the departments and associating the departments with the applications that are used by the departments, you can view the information that is collected about the departments on the Departments page and on the details page for the department.

What to do next

You want to target what department and what storage resources might be responsible for a change or an upward trend in significant storage usage. By correctly classifying the department or departments you can properly plan to archive or switch the storage within the organization.

Related concepts

- [Departments](#)

Related tasks

- [Creating departments](#)
- [Adding and creating subdepartments](#)
- [Viewing information about departments](#)
- [Creating applications](#)
- [Adding subcomponents](#)

Using applications and subcomponents to monitor capacity and space usage

To monitor the performance, capacity, and space usage of the applications in your business organization, create applications and subcomponents. You can also add applications to departments so that storage capacity and usage can be monitored in an overall business hierarchical manner.

Before you begin

You can create a simple application model, which consists of an application and use a filter to associate the storage resources that the application uses with the application. Alternatively, you can create a complex application model that comprises an application and application subcomponents and then add filters to associate the storage resources that each application subcomponent uses with each of the subcomponents.

Besides creating applications to monitor the performance, capacity, and space usage of the applications in your business organization, you also create applications to add them to departments. When you add applications to departments, you can monitor the capacity and space usage of the departments and you can monitor the performance of the storage resources that are associated with the applications that each department uses.


To monitor the capacity of an application, you create an application model that is based on the following elements:

- The application
- The subcomponents that the application contains
- The applications, application subcomponents, or both that are added to the department

In the following tutorial, you want to create an application model that monitors the capacity, space usage, and performance of the Book Sales DB application. The application model that you want to create comprises these application subcomponents:

- Wholesale Transactions
- Retail Transactions
- Online Transactions

Figure 1. An application with a three subcomponent hierarchy



To understand how capacity is trending for storage in an application and subcomponents, complete these tasks:

1. Create the Book Sales DB application.
2. Create the Wholesale Transactions, Retail Transactions and Online Transactions applications to represent the subcomponents.
3. Add the Wholesale Transactions, Retail Transactions, and Online Transactions as subcomponents to the Book Sales DB application.
4. Create the Book Sales department and the Wholesale, Retail, and Online subdepartments to be associated with the applications.

- **[Tutorial: Viewing storage capacity and usage trends](#)**

You can view the capacity and usage trends of storage resources used by an application and subcomponents to diagnose how much storage a particular application is using in relation to the other applications and departments in the business hierarchy.

Related concepts

- [Applications](#)

Related tasks

- [Creating applications](#)
- [Adding subcomponents](#)
- [Viewing information about applications](#)

Tutorial: Viewing storage capacity and usage trends

You can view the capacity and usage trends of storage resources used by an application and subcomponents to diagnose how much storage a particular application is using in relation to the other applications and departments in the business hierarchy.

About this task

You want to understand how capacity is trending for storage that is used by the Book Sales DB application and the Wholesale Transaction, Retail Transactions and Online Transactions subcomponents and the association the application has with the Book Sales department and subdepartments.

Procedure

The storage resources that are used by the application and that are used by each subcomponent are associated with the application and its subcomponents. For example, the storage resources are assigned as follows:

1. Create a database application called Book Sales DB and assign resources that use the volumes on the IBM® Storwize® V7000 with names that begin with bksales for storing sales transactions.
2. Create the subcomponent applications named Wholesale Transaction, Retail Transactions and Online Transactions.
3. Assign the volumes on the IBM Storwize V7000 with names that begin with bksales_ws to the Wholesale Transactions subcomponent, which is used for storing wholesale sales transactions.
4. Assign the volumes on the IBM Storwize V7000 with names that begin with bksales_rt to the Retail Transactions subcomponent, which is used for storing retail sales transactions.
5. Assign the volumes on the IBM Storwize V7000 with names that begin with bksales_ol to the Online Transactions subcomponent, which is used for storing online sales transactions.


Note: On the Create Filter page, you can specify a name pattern to determine which volumes to include. For volumes, you can specify name patterns to determine from which servers, storage systems, or pools the volumes are selected. You can then click Preview to view the volumes that are selected for inclusion in your application.

Results

To monitor space usage for the Book Sales department and its subdepartments, you associate the Book Sales DB application and the Wholesale, Retail, and Online Transactions subcomponents of the application as follows:

- You associate the Book Sales DB application with the Book Sales department.
- You associate the Wholesale subdepartment with the Wholesale Transactions subcomponent.
- You associate the Retail subdepartment with the Retail Transactions subcomponent.
- You associate the Online subdepartment with the Online Transactions subcomponent.

Figure 1. The application, subcomponent, assigned resources and department association



What to do next

You can view the capacity information that is collected about the application on the Applications page. You can also view capacity information for each application and subcomponent on the details page. When you complete adding the departments and associating the departments with the applications that are used by the departments, you can view the information that is collected about the departments on the Departments page and on the details page for the department.

Related concepts

- [Applications](#)

Related tasks

- [Creating applications](#)
- [Adding subcomponents](#)
- [Adding applications to departments](#)
- [Viewing information about applications](#)

Tutorial: Viewing NPIV connections between server ports and switch ports in a fabric

In this tutorial, a storage administrator views the N-Port ID Virtualization (NPIV) connections between the server ports and the switch port in the fabric. With NPIV, multiple node ports can be logically connected to one switch port.

About this task

You want to view connectivity information for a switch port that has NPIV connections.

Procedure

1. In the menu bar, go to **Network > Switches**. Double-click the switch that you want to view, or right-click the switch and click **View Details**.
Tip: Alternatively, in the menu bar, go to **Network > Fabrics**. Double-click the fabric that you want to view, or right-click the fabric and click **View Details**. In the **Internal Resources** section, click **Switches**. Double-click the switch that you want to view, or right-click the switch and click **View Details**.
2. In the **Internal Resources** section, click **Ports**. You can view details about the ports, including the number of NPIV connections.
3. In the **Connected NPIV Ports** column for the port that you want to view, click the number of ports. You can view the following information about the NPIV connections:
 - In the **Connected NPIV Ports** tab, you can view details about the NPIV connections, such as the port type and the name of the connected resource. The port type can be **N_Port** (end node port) or **NP_Port** (proxy node port).
 - In the **Connectivity** tab, you can view details about the physically connected end-node port.
Tip: If IBM Spectrum® Control does not differentiate between the physical and logical connections for the switch port, the details about the NPIV connections are displayed in the **Connectivity** tab. The **Connected NPIV Ports** tab is not displayed.

Tutorial: Exporting and uploading performance data for a SAN Volume Controller system

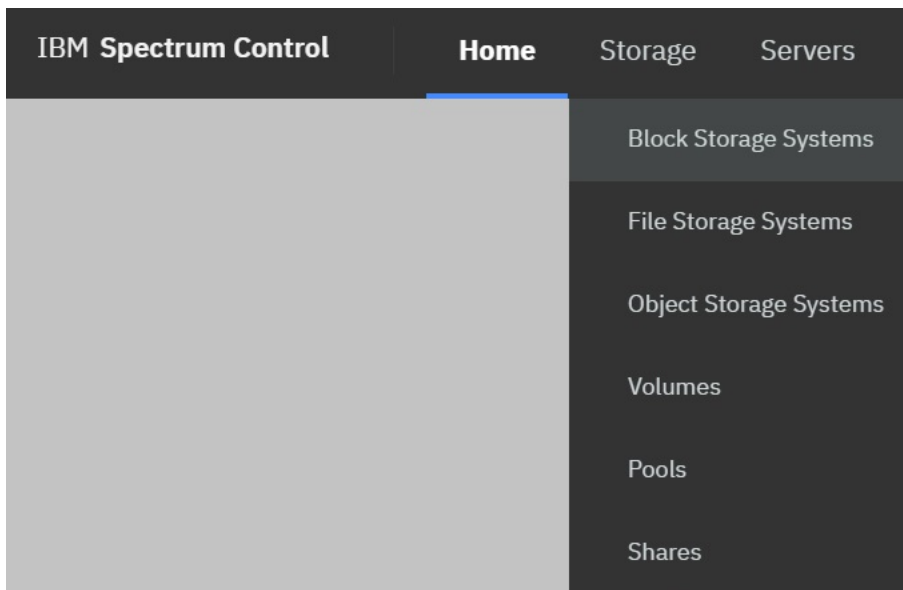
In this tutorial, you export performance data for a SAN Volume Controller to a compressed package » and upload it to a support ticket. « You can also send the package to IBM® Support.

About this task

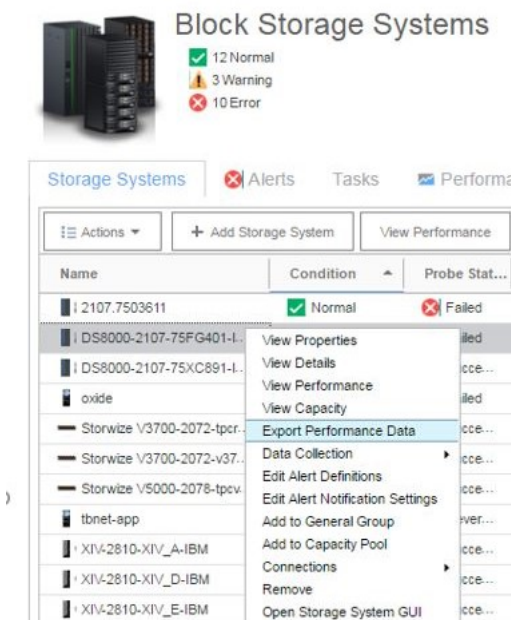
Some volumes on a SAN Volume Controller system have performance problems. You consult IBM Support who require detailed performance data about the SAN Volume Controller system to diagnose the problem. You are asked to export performance data by using IBM Spectrum® Control.

Procedure

1. Go to **Storage > Block Storage Systems**.



2. Right-click the SAN Volume Controller system, then click Export Performance Data.




3. For the Export type, specify whether to download performance data to a local file or to download it to a local file and upload it to a support ticket. For this scenario, select Upload to Ticket.

»You must select the ticket number from the list or manually type it in the Ticket field.

4. Select a time range for the support package.

Typically, the time range includes the time when the performance problem occurred.

5. If you are instructed to do so, select the Advanced export check box. Click Done.



»A task is started and shown in the Running and Completed tasks icon on the menu bar.

6. »When the task is complete, click the Download icon in the task to save the file locally.

Results

»The package is downloaded and uploaded to a support ticket.

Related reference

- [Exporting performance data for storage systems and fabrics](#)

Tutorial: Comparing the performance of storage systems

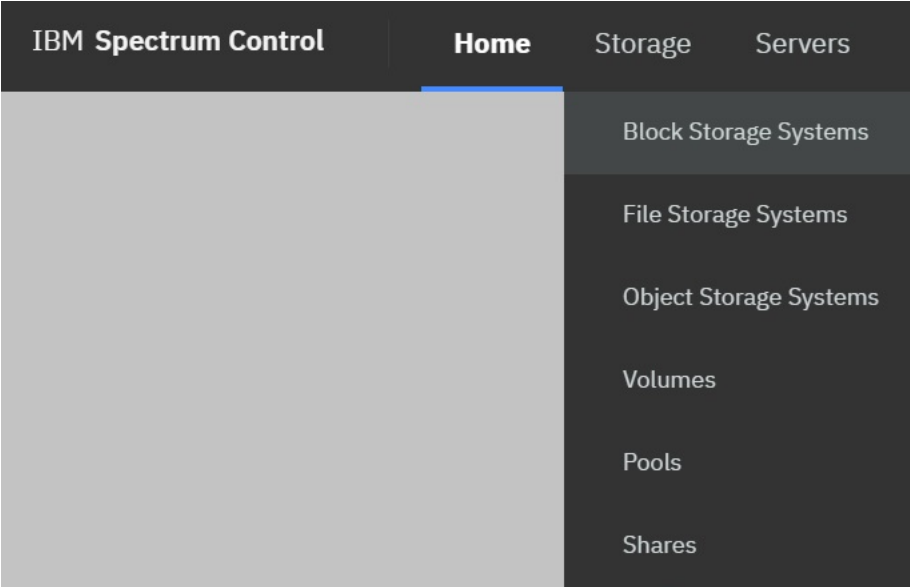
In this tutorial, you compare the performance of two storage systems.

About this task

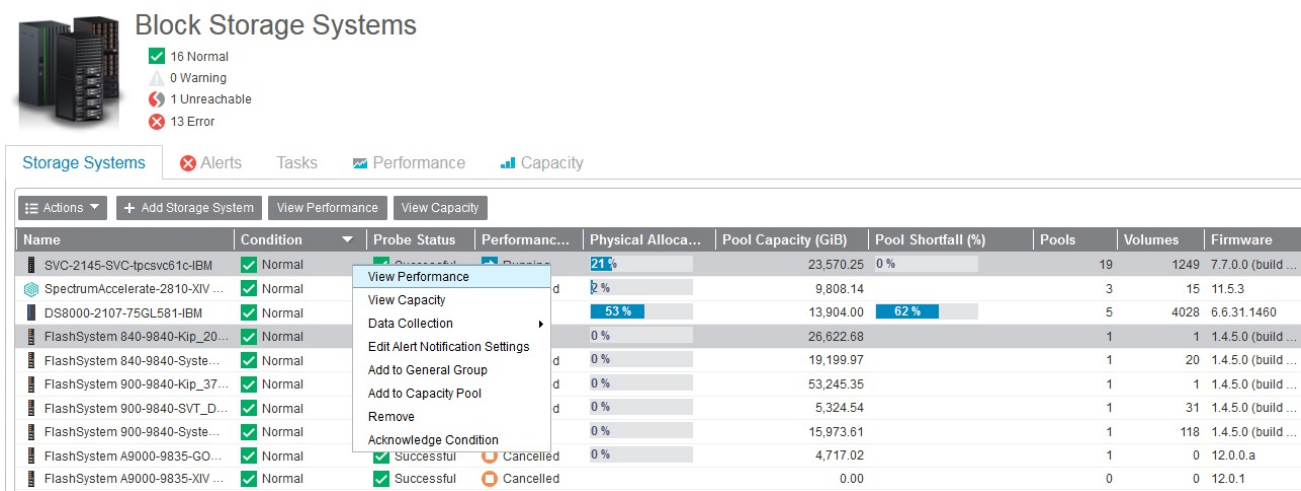
Your company recently purchased an IBM FlashSystem® 900 storage system. The storage system was added to IBM Spectrum® Control and the performance of the storage system is being monitored. Your managers ask you to compare the performance of the FlashSystem 900 storage system to another storage system so that they can assess the performance of their investment. In particular, your managers are interested in the response times and the I/O rates of the FlashSystem 900.

Procedure

1. Go to Storage > Block Storage Systems.




2. Click the FlashSystem 900 storage system and the non-flash storage system that you want to compare the FlashSystem 900 with.
3. Right-click the selected storage systems, then click View Performance.



Results

A performance view is displayed. The view shows the performance of the two storage systems on a chart and in a table. By default, the performance view shows the Total I/O Rate - overall and Overall Response Time metrics. The time period for the information is the last 12 hours.



Tutorial: Reviewing and updating your agentless servers

In this tutorial, you review the agentless servers that are created automatically by IBM Spectrum® Control, and update your agentless servers.

About this task

Every month, you create a chargeback report for the physical servers in your storage environment. The report shows how much capacity is used by physical servers, and the cost of that capacity.

IBM Spectrum Control creates and updates agentless servers automatically after storage systems and hypervisors are probed. The agentless server information is used to create an accurate model of your storage environment, which is then used by the chargeback report to show the consumption of storage by server. You want to review your agentless servers to make sure that they accurately represent your storage environment.

Procedure

1. Go to Servers > Servers.
2. Review the servers that have Agentless in the Condition column.
You notice that IBM Spectrum Control identified 4 host connections and created an agentless server for each host connection. However, in your environment, these 4 host connections are on one server. You need to combine the 4 agentless servers into a server.
3. Select the four agentless servers, right-click, and then click Merge Servers.

Servers

3 Normal
0 Warning
0 Error
1031 Agentless

Servers Alerts Tasks Capacity

Actions + Deploy Agent View Capacity

Name	Con...	Probe Sta...	Agent State	OS Type
basalt.storage.tucso...	Agent...	N/A		Windows
cet-hacmp1.storage....	Agent...	N/A		Other
cet-hacmp2.storage....	Agent...	N/A		Other
cet_hacmp5	Agent...	N/A		Other
cet_hacmp6	Agent...	N/A		Other
Charge_zLargeEn...				Other
ChostWWPN				Linux
Consumer_zLarge...				Other
crawfish.storage.tu...				Other
dummy_test_host...				Other
EmptyHost41				Windows
ESX_abbott	Agent...	N/A		Other

4. Type a name for the new server, then click Merge.

5. Review the agentless servers again.

You notice that IBM Spectrum Control created an agentless server from a group of host connections that appear to be related. However, in your environment, these host connections are actually individual servers. You need to separate the agentless server into individual servers.

6. Select the agentless server, right-click, and then click Separate Server.

Servers

3 Normal
0 Warning
0 Error
1031 Agentless

Servers Alerts Tasks Capacity

Actions + Deploy Agent View Capacity

Name	Con...	Probe Sta...	Agent State	OS Type
ESX_abbott	Agent...	N/A		Other
excitebike.storage.tu...	Agent...	N/A		Other
farro1.storage.tucso...				Other
fern.storage.tucson.i...				Other
flint_2				Other
galati.storage.tucson...				AIX
hops2.storage.tucso...				Windows
HostEmptyTest				AIX
IFS-Ballistic				Other
IFS-tpcrifs1_bellagio				Other
killerbee.storage.tuc...				Linux
maize2.storage.tucs...				AIX
mdm-c27_RENAMED	Agent...	N/A		Other

7. If any WWPNs on the server are used by more than one host connection, an error message is displayed. The message displays the WWPNs that are used by more than one host connection. Review the WWPNs and make sure that each WWPN is used by only one host connection. Then, try to separate the server again.

8. On the Separate Servers dialog, click Apply.

On the resource list page for servers, an agentless server is created for each of the servers.

Related concepts

- [Agentless servers](#)

Tutorial: Troubleshooting performance

Bob is a storage administrator. One morning, he receives a ticket: A critical medical application, Epic Database, has a performance problem. Bob must investigate the problem and identify the cause.

About this task

Previously, Bob created an Epic Database application on IBM Spectrum® Control and specified the resources that the application uses.

Procedure

1. Bob clicks Applications on the dashboard to troubleshoot only the Epic Database resources.

[Applications](#) [17](#)




2. Bob right-clicks the Epic Database application on the Applications page, then clicks View Details.



Applications

Applications						
<div> <div> <div>⋮</div> <div>Actions</div> </div> <div> <div>+</div> <div>Create Application</div> </div> </div>						
Name	Block Capacity (GiB)	File Capacity (GiB)	Volumes	Servers	Hypervisors	
Accelerate Group	112.18	0.00	7			
Cigna Online	64.11	0.00	4			
Citrix	528.00	0.00	34	tivfcm		
Connections Nodes	112.00	0.00	6			
DB2 Database Ser...	450.00	0.00	9	tpcaix01.democent...		
demo database	416.00	0.00	28	tivfcm		
DemoCentral vm...	205,570.21	0.00	101			
Epic Database	568.00	0.00	12			
Epic Reporting		0.00	9			
Exchange		0.00	29			
Exchange Mail Box...		0.00		2	2	
FileSet SPinst1		160.00				
Intranet Web Servers		0.00	12			

Bob immediately notices the high response times of one volume, vsc-epic-db-test-goldp1, in the Most Active Volumes chart. He sees at a glance that this volume is well above 10 ms/op.



Bob is familiar with the response times and other network conditions that are acceptable for the application. He knows that this response time is too high for this application. He also notes that the high response time is sustained, not just spikes. Bob decides to look more closely at the performance of the volume.

3. Bob clicks Volumes in the Related Resources section, then clicks the Performance tab.

4. Bob uses the chart and table in the Performance tab to further investigate the cause of the problem with the Epic Database application.

He uses the chart and table in the Performance tab to investigate by using the following techniques:

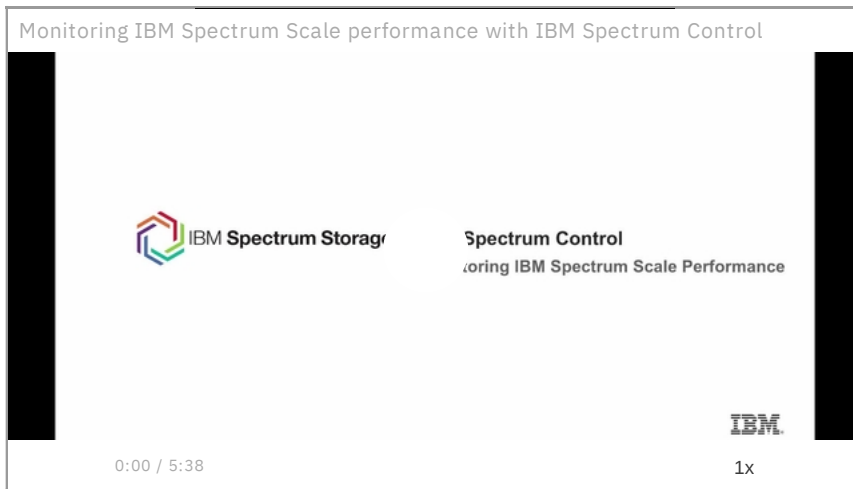
- Change which metrics are displayed in the performance charts. He clicks a metric in the chart legend to display only that metric. He also clicks the Select Chart Metrics icon to add other metrics to the chart.
- Change which resources are displayed in the performance charts. He clicks a resource in the legend to display only that metric. He also clicks resources in the table to add them to the chart.
- View the relationships between resources.
- Compare metrics by time range. To make this comparison, he can select either a predefined time range or a custom time range.
- Compare different metrics within the same time range.
- Compare the same metrics across different time ranges.
- Identify the best or worst values for a performance metric from a large list of resources. For example, Bob might want to see which of the managed disks is handling the most workload, or which volumes have the highest response times.
- Compare performance charts side by side. To compare charts, Bob opens a chart in a new window and arranges the charts beside one another.
- Synchronize all charts so that they all show data for the same time period.
- Export the data from a chart in CSV format.

Tutorial: Monitoring IBM Spectrum Scale performance

You can use IBM Spectrum® Control to monitor the performance of the IBM Spectrum Scale clusters in your storage environment.

In the following video, you learn how to use the following techniques to identify performance issues and bottlenecks in your storage:

- Analyze the performance of all your clusters over a period to determine which clusters have the heaviest I/O loads or the longest response times.
- View which of the nodes that mount a file system have the heaviest I/O loads and highest CPU utilization.
- Compare the file systems on a cluster to determine which file systems have the heaviest I/O loads and the longest response times.
- Monitor the performance of SAN attached volumes that are backing the Network Shared Disks (NSDs) in the cluster. If you have block storage devices, such as IBM Spectrum Accelerate or IBM FlashSystem®, that are providing storage for your clusters, you can monitor the block volumes that are backing the NSDs.



Tutorial: Viewing the aggregated workload for an application

Bob is a storage administrator. Bob wants to easily evaluate the performance of the accounting application because the Accounts department is expanding to a new location and they want to replicate their data across both locations. Bob needs to work out the inter-site link capacity requirements.

About this task

Previously, Bob created an accounting application, Accounts Database, and specified the resources that the application uses.

Procedure

1. Bob clicks Applications on the dashboard to view only the Accounts Database resources.

Applications 17



2. Bob right-clicks the Accounts Db application on the Applications page, then clicks View Details.

Applications

Applications Capacity


Actions Create Application Create Report View Performance View Capacity

Name	Block Capacity (GiB)	File Capacity (GiB)	Object Capacity (GiB)	Volumes	Servers	Hypervisors
Accounts DB	452.48	0.00	0.00	52	4	

View Properties View Details View Performance View Alert Definitions View Capacity Add to... Delete


Bob can use the Aggregated Workload chart to quickly evaluate the throughput required for the link between both sites.

Overview



Bob can toggle between seeing the aggregated I/O rate and the data rate for the application using the arrow (↔) icons either side of the chart.

Overview



3. The rates are aggregated on the primary volumes associated with the application which allows Bob to accurately assess the required MiB/s throughput he will need between both locations.

Tutorial: Identifying the source of slow drain problems caused by depletion of buffer credits

Use this tutorial to find out how to use IBM Spectrum® Control to identify a host that has depleted buffer credits that are causing a slow drain condition.

About this task


Fibre Channel (FC) networks use buffer credits to control the flow of data frames from port to port. The number of buffer credits for a port is the number of data frames that the port can receive. When that number is reached, ports cannot send further data frames until the receiving port indicates that it is ready. If all of the buffer credits of a port are being used, then the port cannot receive more data.

For example, if a host has a performance problem, then its ports might not be able to clear their buffer credits to receive more data. If the host ports cannot receive data, then switch ports cannot send data to the host ports, so the buffer credits of the switch port become depleted too. Ports on other switches in the fabric that try to send data through the switch port are also affected, and their buffer credits become depleted in turn. In this way, the buffer credit problem builds throughout the storage environment. The buffer credit depletion on the host ports impacts the switches that communicated with the host. The switches that communicate with that switch cannot use their buffer credits, so storage systems cannot communicate with the switches.

In this way, a single host with a performance problem can impact all the hosts that use the same switches and inter-switch links. This condition is called *slow drain*. Slow drain in your storage environment can manifest as a problem with storage systems rather than with a host.

Procedure

1. Configure a performance alert for Port Send Delay Time, Port Send Delay I/O Percentage, or Zero Buffer Credit Timer, depending on the storage system. For example, configure a Port Send Delay I/O Percentage alert for FlashSystem 9100 that is triggered when the delay is greater than 20%.



FlashSystem-AF7-9100-1
IBM FlashSystem 9100 - 9848

Actions

General

- Overview
- Performance
- Properties
- Alerts (59)**
- Alert Definitions
- Tasks (0)
- Data Collection (2)

Internal Resources

- Volumes (32)
- Pools (1)
- Managed Disks (1)
- RAID Arrays (1)
- Drives (12)
- I/O Groups (1)
- Nodes (2)
- Enclosures (1)
- FC Ports (16)
- Host Connections (3)

Related Resources

- Servers (1)
- Fabrics (5)
- Switches (5)

Alerts

0 Critical
47 Warning
12 Informational

Alert Policy: None Policy Actions


Actions Refresh

Alert Name Condition Violation Sev... Occurrence Time Internal Resource Alert Source

IBMPA_PerformancePortSendDela...	>= 20 %	33.35 %	Warn...	Apr 7, 2020, 19:28:20	node1.Port 3	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	33.43 %	Warn...	Apr 7, 2020, 19:28:20	node1.Port 1	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	31.95 %	Warn...	Apr 7, 2020, 19:28:20	node1.Port 4	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	32.58 %	Warn...	Apr 7, 2020, 19:28:20	node1.Port 2	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	22.94 %	Warn...	Apr 7, 2020, 23:18:23	node1.Port 4	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	24.26 %	Warn...	Apr 7, 2020, 23:18:23	node1.Port 3	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	24.35 %	Warn...	Apr 7, 2020, 23:18:23	node1.Port 1	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	23.54 %	Warn...	Apr 7, 2020, 23:18:23	node1.Port 2	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	31.73 %	Warn...	Apr 8, 2020, 00:49:21	node1.Port 4	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	32.38 %	Warn...	Apr 8, 2020, 00:49:21	node1.Port 2	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	33.32 %	Warn...	Apr 8, 2020, 00:49:21	node1.Port 1	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	33.32 %	Warn...	Apr 8, 2020, 00:49:21	node1.Port 3	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	29.13 %	Warn...	Apr 8, 2020, 01:59:24	node1.Port 3	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	29.04 %	Warn...	Apr 8, 2020, 01:59:24	node1.Port 1	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	27.64 %	Warn...	Apr 8, 2020, 01:59:24	node1.Port 4	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	27.99 %	Warn...	Apr 8, 2020, 01:59:24	node1.Port 2	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	28.99 %	Warn...	Apr 8, 2020, 03:14:20	node1.Port 2	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	30.34 %	Warn...	Apr 8, 2020, 03:14:20	node1.Port 3	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	30.44 %	Warn...	Apr 8, 2020, 03:14:20	node1.Port 1	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	28.87 %	Warn...	Apr 8, 2020, 03:14:20	node1.Port 4	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	33.47 %	Warn...	Apr 8, 2020, 03:34:21	node1.Port 4	Default FlashSystem Family policy
IBMPA_PerformancePortSendDela...	>= 20 %	34.95 %	Warn...	Apr 8, 2020, 03:34:21	node1.Port 3	Default FlashSystem Family policy

Showing 59 items | Selected 0 items Refreshed a few moments ago

- To view the alerts, click Home and then click Alerts. If the Port Send Delay Time alert was triggered, note the time of the alert.
- To view information about the affected storage system, click the link in the Resource column.
- In the Internal Resources section of the storage system details page, click Volumes.
- Click the Performance tab.
- Set a time period for the performance chart. Set the start time to before the alert occurred and the end time to after the alert occurred.
- Set the chart to display the following metrics:
 - Read Data Rate
 - Overall Response time
- Sort the performance table by the Max Total I/O Rate column.
- Click the volume with the highest Total I/O Rate to show the volume in the chart. Verify that the Read Data Rate spiked at the time that the alert occurred.



FlashSystem-AF7-9100-1
IBM FlashSystem 9100 - 9848

Actions

General

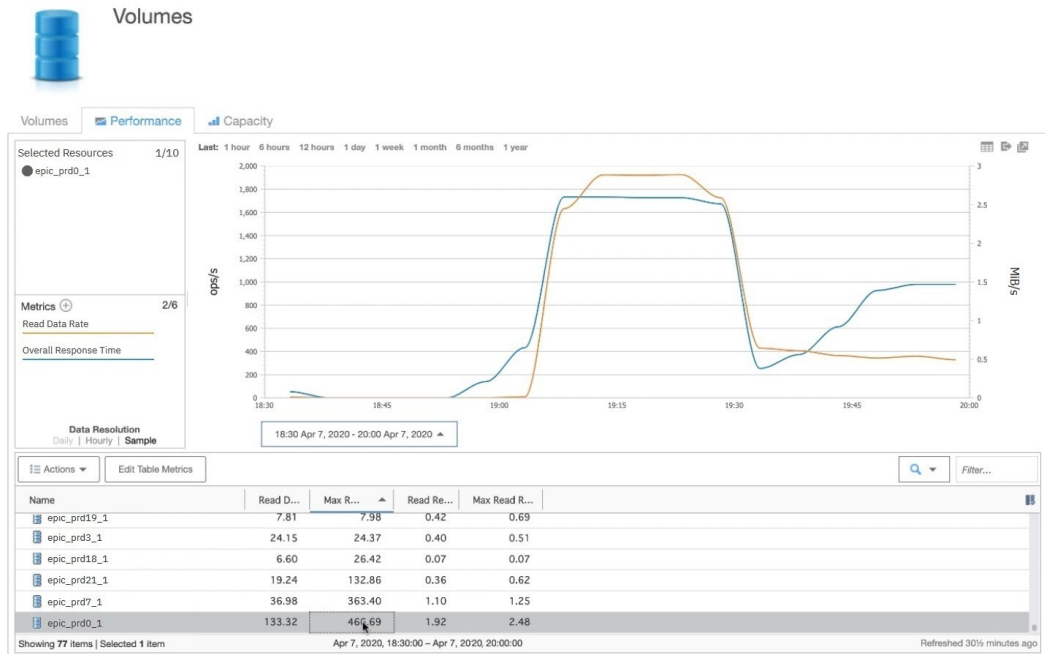
- Overview
- Performance
- Properties
- Alerts (59)**
- Alert Definitions
- Tasks (0)
- Data Collection (2)

Internal Resources


- Volumes (32)
- Pools (1)
- Managed Disks (1)
- RAID Arrays (1)
- Drives (12)
- I/O Groups (1)
- Nodes (2)
- Enclosures (1)
- FC Ports (16)
- Host Connections (3)

Related Resources

- Servers (1)
- Fabrics (5)
- Switches (5)



- Right-click the volume in the performance table, then click Host Connection Performance.
- Set the chart to display the following metrics:
 - Read Data Rate
 - Overall Response time
- Verify that the spike for the Read Data Rate and the Overall Response Time occurred when the alert was triggered. If the spike occurred at the same time, the host **P9_lpar5_epictest01** is the source of the slow drain that is causing problems on the storage system.



Tip: If the buffer credit problem occurs in a cluster, there might be multiple hosts that are mapped to the volume. In this case, you must investigate the hosts individually to determine which ones cause the delay in sending from the port.

Tutorial: Identifying the locations of devices

Identify the locations of devices that are monitored by IBM Spectrum® Control. You can also apply custom tags to more easily identify, sort, or group devices based on location or another attribute that you specify.

About this task

In large or distributed environments, it can be a challenge to manually keep track of all the physical locations where your devices reside. This challenge might be especially true in organizations with data centers that span multiple time zones and countries. In IBM Spectrum Control, when you add a device for monitoring, it automatically tracks the locations and time zones of devices for you.

Use the following steps to identify the location and time zone of devices that are being monitored. You can also apply custom tags to better filter or sort the devices in the GUI or in an external application if the data is shared or exported.

Procedure

1. In the menu bar, select the device type.
For example, if you want to identify the locations of block storage systems, go to **Storage > Block Storage Systems**.
2. Right-click a device and select **View Properties**.
3. On the **General** tab, scroll down until you see the **Time Zone** and **Location** values. Use these values to identify where a device resides.
To further refine location information, such as differentiating devices in the same geographic location but separate data centers, you can apply up to three custom tags.
4. Optional: To apply a custom tag to a device, click **Edit**.
5. Optional: Enter a value in a **Custom Tag** field and click **Save**.
Use custom tags to help you more easily identify, sort, or group devices. For example, you might enter values like **Data Center X** or **Data Center Y**, depending on the location of a device.
Tip: You can bulk apply custom tags to multiple devices at the same time. For example, on the **Block Storage Systems** page, click **Ctrl + click** or **Shift + click** to select multiple storage systems. Then, right-click the selected storage systems, click **Edit Properties**, and enter values in the **Custom Tag** fields.

Example

After custom tags are applied to devices, you can sort, filter, group, and export the devices based on those tags.

For example, on the **Block Storage Systems** page, you can sort the list so that storage systems with a specific custom tag, such as **Data Center X**, are shown at the top of the list. To sort the list, right-click any of the column headings and select the check box next to **Custom Tag 1**. Then, click the **Custom Tag 1** column to sort the list.

You can also filter the list to only show storage systems with a specific custom tag. To filter the list, enter that custom tag value in the **Filter** text box and press **Enter**. Only the storage systems with custom tags that match or partially match the filter text are shown on the page. You can also complete the following actions:

- To group the filtered storage systems, click **Ctrl + click** or **Shift + click** to select them. Then, right-click the selected storage systems, click **Add to General Group**, and follow the prompts.
- To export the filtered storage systems to a file, select **Actions > Export** and select the file format that you want to use.

Planning

The following sections provide information to help plan your IBM Spectrum® Control environment before you install the product. In a complex environment, good planning helps you avoid delays and problems in getting your system up and running.

- **Planning for installation**
You can install IBM Spectrum Control by using the installation program or the command line in silent mode. Installing IBM Spectrum Control by using console mode is not supported. Also, upgrading the operating system of a server where IBM Spectrum Control is installed is not supported. If you must upgrade the server, contact the IBM® Support team for IBM Spectrum Control Support to discuss your options.

- [Planning for configuration](#)
Use this information to plan your IBM Spectrum Control environment. The information includes general configuration guidelines, TCP/IP ports, user names and user rights, and SMI-S support.
- [Planning for storage management](#)
Plan for how to use IBM Spectrum Control to manage the storage and storage resources in your environment.
- [Sudo command privileges](#)
When you install, operate, or uninstall IBM Spectrum Control as a non-root user on AIX® or Linux® operating systems, you are required to have sudo privileges to multiple commands. Use command aliases in the `/etc/sudoers` file to configure your sudo privileges for multiple commands.

Planning for installation

You can install IBM Spectrum® Control by using the installation program or the command line in silent mode. Installing IBM Spectrum Control by using console mode is not supported. Also, upgrading the operating system of a server where IBM Spectrum Control is installed is not supported. If you must upgrade the server, contact the IBM® Support team for IBM Spectrum Control Support to discuss your options.

Antivirus software considerations

Installing IBM Spectrum Control involves making use of your operating system in manners typical for installing new application software. If your antivirus software is set on the maximum mode, it might prevent some of your changes from being accepted.

To verify that your installation completes correctly, enable your antivirus software product to allow the following instances:

- The `/etc/hosts` file can be edited
- Files can be created in the `/temp` directory
- New executable files can be created in the `C:\Program Files` directory

McAfee tip for Windows: If McAfee Adaptive Threat Protection is enabled on the server where IBM Spectrum Control is installed, it might prevent some services from starting or stopping. To help avoid this issue, open McAfee Adaptive Threat Protection and go to settings. In the Real Protect Scanning (Windows only) section, verify if Enable client-based scanning is selected. If so, select Low from the Sensitivity level list.

For more information about McAfee Adaptive Threat Protection, see the following links:

- [Adaptive Threat Protection — Options](#)
- [Overview of Adaptive Threat Protection](#)

Recommended skills and knowledge

Administrators and users of IBM Spectrum Control should be familiar with the following skills and topics:

- General procedures for installing software on Microsoft Windows, IBM AIX®, and Linux® operating systems
- Storage Area Network (SAN) concepts
- IBM Spectrum Control concepts
- IBM Database
- Simple Network Management Protocol (SNMP) concepts

Required user privileges for installation

Ensure that you have the required privileges to install Db2® and IBM Spectrum Control. If you are installing on a Windows operating system, you must install as a user with administrator privileges. If you are installing on an AIX or Linux operating system, you must install as a root user or as a non-root user who has sudo privileges.

Fully qualified host names

Some systems might be configured to return a short host name such as `server22`, instead of a fully qualified host name such as `server22.myorg.mycompany.com`. IBM Spectrum Control requires fully qualified host names, so you must install the product on a computer that has a fully qualified host name.

Restriction:

Hostnames must be 15 characters or less and cannot contain special characters, for example, a space or an underscore. Use characters a through z, or characters 0 - 9. Hyphens (-) are also allowed.

Important: Before you run your Db2 installation and begin your IBM Spectrum Control installation, please validate that the Windows server hostname is 15 characters or less in length.

For more information about determining whether you have a fully qualified domain name, see [Checking for a fully qualified host name](#).

Self signed certificate considerations

There are new certificate requirements from the [CA/B Forum](#) that are strictly enforced for macOS Catalina users that might affect your ability to access the IBM Spectrum Control GUI. During an upgrade of IBM Spectrum Control, certificates self-signed by IBM Spectrum Control will be made automatically compliant. However, if one or more of your certificates are not self-signed by IBM Spectrum Control, see [IBM Spectrum Control and macOS Catalina \(10.15\) Increased Security Policies on SSL Certificates](#) and validate that your certificates are compliant.

Starting with IBM Spectrum Control 5.3.6, all default self-signed certificates have an expiration date of 825 days based on the new certificate requirements from the CA/B Forum. By default, upgrades of IBM Spectrum Control always renew the expiration of all default self-signed certificates used by the product for another 825 days. However, if you have installed your own certificates, IBM Spectrum Control does not modify those certificates for you.

For more information about determining whether you have a fully qualified domain name, see [Replacing the default SSL certificate for the Device, Alert, or Web server with a self-signed certificate](#).

For more information about determining whether you have a fully qualified domain name, see [Replacing the default SSL certificate for the Export server](#).

Hardware and software requirements

The IBM Spectrum Control server can require a large amount of memory, disk space, network bandwidth, and processor resources. To promote consistent availability and performance, install IBM Spectrum Control on a dedicated server where no other critical business applications, including IBM Copy Services Manager, are installed. IBM Software Support is not responsible for availability, performance, and functional problems that are caused by other business applications that are deployed on the same server as IBM Spectrum Control.

Before you install IBM Spectrum Control, apply all available updates and patches to the operating system on the target computer.

For information about hardware requirements, see [Hardware requirements](#).

For information about software requirements, see [Software requirements](#).

For information about updates to product and operating system support, see *Find the Supported Hardware, Products and Platforms Interoperability Matrix Links* at <https://www.ibm.com/support/pages/node/388393>.

Installation methods

You can install Db2 and IBM Spectrum Control by using the GUI or silent mode. In silent mode, a command is provided with the values in a response file.

Use the silent-mode installation to install on a server that cannot display graphics. For example, if you are installing on an AIX or Linux server that does not have the X Window System installed.

On a server that does not have the X Window System installed, you cannot change the Db2 and IBM Spectrum Control passwords within IBM Spectrum Control by using the change password tool. For information on how to change the passwords, see [Changing passwords on AIX and Linux systems using the Command Line Interface \(CLI\)](#).

Installation images

Compressed installation images are available electronically. Download and extract these images to a location with adequate disk space. For disk space requirements by operating system, see *IBM Spectrum Control - Hardware Support: Memory, Processor and Disk Space* at [IBM Spectrum Control - Hardware Support: Memory, Processor and Disk Space](#).

Download the installation images for each product into a separate directory. The following list provides sample download directory paths:

- c:\downloads\DB2
- c:\downloads\SC

Extract the image files into a separate temporary directory for each product. For example:

- c:\temp\DB2
- c:\temp\SC

Attention: If you do not extract the files into separate directories for each product, the common paths or path structures might overwrite each other, and you cannot install the products.

Follow these rules for naming the directories:

- Do not include spaces in the directory name.
- Keep the directory name short.
- For Linux or AIX operating systems, do not include a period (.) at the end of the directory name.

During the extraction of multi-part images, if you are prompted to merge or replace folders or files with the same name, click Merge/Replace.

Installation files for the following components are provided with IBM Spectrum Control:

Db2

- Db2 for Windows (64-bit)
- Db2 for Linux (64-bit)
- Db2 for AIX (64-bit)

IBM Spectrum Control

- IBM Spectrum Control for AIX
- IBM Spectrum Control for Linux
- IBM Spectrum Control for Windows

In addition, installation files for the following optional agents are available with IBM Spectrum Control:

- IBM Spectrum Control Storage Resource agent for Windows
- IBM Spectrum Control Storage Resource agent for UNIX

On Windows operating systems, assign the appropriate security settings for the Db2 and IBM Spectrum Control compressed files by completing the following steps:

1. Right-click one of the compressed installation files.
2. Examine the General tab and determine whether the Unblock button is displayed.
3. If the Unblock button is displayed:
 - a. Right-click the compressed file and click Properties.
 - b. On the General tab, next to Security, click Unblock.
4. Verify that the compressed installation files no longer have the Unblock button displayed.

IBM Spectrum Control GUI

After IBM Spectrum Control is installed, use the GUI to complete your storage management tasks. The GUI runs in a web browser and includes the ability to monitor, manage, and troubleshoot storage resources. You can access this interface from anywhere that you have a web browser and connectivity to a network. You can also use the IBM Cognos® Analytics reporting tool to generate, save, and view reports about the condition, capacity, and performance resources.

Installation language

When you install IBM Spectrum Control, you must select one of the following languages:

- English
- Czech
- French
- German
- Hungarian
- Italian
- Japanese
- Korean
- Polish
- Brazilian Portuguese
- Russian
- Spanish
- Chinese (Simplified)
- Chinese (Traditional)

For more information about changing the operating system for the user interface, see [Changing languages](#).

Messages, online help, and text are displayed in the language that you select. You are not required to install a language pack after installation. When you use IBM Spectrum Control, the language that is displayed is the language setting of the operating system.

To install IBM Spectrum Control on a Windows domain, see [Planning to install IBM Spectrum Control in a Windows domain](#) and [Installing IBM Spectrum Control on a Windows domain](#).

Installation images on AIX

When you extract installation images, use the GNU tape archive (tar) program instead of the AIX tape archive program. The AIX tape archive program might truncate long file names, which can cause installation errors in IBM Spectrum Control.

To get the GNU tape archive program, go to [AIX Toolbox for Open Source Software](#). Find tar in the Package column, and click RPM to download the program. You must specify this program as the default tape archive program in the PATH environment variable.

For more information on using the correct program to extract the IBM Spectrum Control installation media, see [How to install GNU tar on AIX operating system](#).

- [Planning to install IBM Spectrum Control in a Windows domain](#)
Before you can install IBM Spectrum Control on a Windows domain, you must determine which installation method is appropriate, based on your environment.
- [Required user privileges in installation scenarios](#)
When you are planning to install IBM Spectrum Control, there are users to consider for the installation.

Planning to install IBM Spectrum Control in a Windows domain

Before you can install IBM Spectrum® Control on a Windows domain, you must determine which installation method is appropriate, based on your environment.

Installation methods

Install Db2® and IBM Spectrum Control in one of the following ways:

- Use local user accounts to install both Db2 and IBM Spectrum Control.
- Use a local user account to install Db2 and a domain user account to install IBM Spectrum Control.
- Use domain user accounts to install both Db2 and IBM Spectrum Control.
- [Windows domain and local user accounts](#)
When a computer is a member of a Windows domain, you can install IBM Db2 on the local computer or on a computer that is a member of a Windows domain. The installation process creates a local Db2 user account or a domain Db2 user account.
- [Adding a computer to the Windows domain](#)
Before you can install IBM Spectrum Control in a Windows domain, you must first add the computer on which you plan to install IBM Spectrum Control to the domain.
- [Verifying that the Netlogon service is running](#)
IBM Spectrum Control uses WebSphere® Application Server Liberty to authenticate domain users. WebSphere Application Server Liberty requires that the Microsoft Netlogon service is enabled and running to authenticate these users.
- [Installing Db2 by using a Windows domain user account](#)
Before you install IBM Spectrum Control in a Windows domain, you must install Db2 and register the Db2 license key.
- [Creating a Windows domain common user account for IBM Spectrum Control](#)
You must create a Windows domain common user account before you can install IBM Spectrum Control in a Windows domain.
- [Granting Db2 SYSADM authority to a Windows domain user account](#)
If a Windows domain user account is used to install IBM Spectrum Control, the user account may not have the Db2 SYSADM authority, because Db2 goes to the domain controller computer to list the groups. Before you install IBM Spectrum Control, you must grant the Windows domain user accounts the Db2 SYSADM authority.

Windows domain and local user accounts

When a computer is a member of a Windows domain, you can install IBM® Db2® on the local computer or on a computer that is a member of a Windows domain. The installation process creates a local Db2 user account or a domain Db2 user account.

User accounts

Windows domain user accounts are used to manage multiple computers, and local user accounts can be used to manage one computer.

Important: You set the common user for IBM Spectrum® Control during the initial installation of the product. When that common user is set, it cannot be changed during the life of that IBM Spectrum Control system. For example, if you set the common user to be *db2admin* during the initial installation of IBM Spectrum Control on a Windows operating system, then the *db2admin* user is the constant common user for that particular IBM Spectrum Control system.

Restriction: Before you install IBM Spectrum Control by using a Windows domain or a local user account as the common user name, you must add the Windows domain user account or local user account to the local administrators group. The IBM Spectrum Control installation software recognizes only Windows domain or local user accounts (and not Windows domain or local groups) that are added to the local administrators group.

domain Db2 user account is an example of a Windows domain user account, and *local Db2 user account* is an example of a local user account.

Naming conventions for user accounts

You can install IBM Spectrum Control by using the same user account that is defined in both the local and Windows domain registries. The Windows operating system resolves this naming collision on computers that belong to a Windows domain by prefixing the user name with the host name or domain name. This user name is also called a *fully qualified user name*.

You must use the following naming conventions for each user account:

- *domain_name\administrator*
- *domain_name\user 1, domain_name\user2, domain_name\userN.*
- *host_name\administrator*
- *host_name\user1, host_name\user2, host_name\userN.*

For example, if the domain name is **TPC51**, and the user name is **db2admin**, the Windows domain user account is **TPC51\db2admin**. The naming convention for a local user account is *host_name\user name*. For example, if the host name is **machine2**, and the user name is **db2admin**, the user account is **machine2\db2admin**. You must enter the fully qualified user names in the common user name field when you install IBM Spectrum Control on computers that are members of a Windows domain.

Installation considerations

As part of your planning process, you must determine which user accounts to use when you install IBM Spectrum Control and Db2, consider the environment and security requirements.

IBM Spectrum Control uses WebSphere® Application Server Liberty to authenticate the users for local and domain user accounts. When the computer that hosts the WebSphere Application Server Liberty process is a member of a Windows domain, by default local and domain user registries are used, but the Windows domain user registry takes precedence.

Related reference

- [Adding a computer to the Windows domain](#)
- [Installing Db2 by using a Windows domain user account](#)
- [Creating a Windows domain common user account for IBM Spectrum Control](#)
- [Granting Db2 SYSADM authority to a Windows domain user account](#)
- [Installing IBM Spectrum Control on a Windows domain](#)

Adding a computer to the Windows domain

Before you can install IBM Spectrum® Control in a Windows domain, you must first add the computer on which you plan to install IBM Spectrum Control to the domain.

You can log on to a computer that is a member of a Windows domain by entering a user name that is prefixed with the domain name. For example, you can enter **TPC52\Administrator** if you are the domain administrator of the **TPC52** domain.

To add a computer to a Windows domain, complete the following steps:

1. Click Start > Control Panel > Network and Internet > Network and Sharing Center.
2. On the Network and Sharing Center window, click the left mouse button on your local area connection and click Properties.
3. Select the Internet Protocol Version 4 (TCP/IP4) check box, and click Properties.
Tip: If you are using Windows Server 2012, ensure that the Client for Microsoft Networks check box is selected.
4. On the General tab, select Use the following DNS server addresses.
5. In the Preferred DNS server field, enter the IP address of the domain controller computer.
6. In the Alternate DNS server field, enter the IP address of the alternate DNS server.
7. Click Advanced.
8. On the DNS tab, review the list of IP addresses.
9. In DNS suffix for this connection field, enter the fully qualified domain name.
10. Click OK.
11. Right-click Computer and select Properties.
12. In the Computer name, domain, and workgroup settings section, click Change settings.
13. Click Change.
14. On the Computer Name/Domain Changes window, select Domain, enter the fully qualified domain name, and click OK.
15. Enter the user name and password for the domain controller computer.

- Important: The user name and password that you enter must be a domain administrator user name and password.
- Click OK for the changes to take effect and to start the computer again.

Verifying that the Netlogon service is running

IBM Spectrum® Control uses WebSphere® Application Server Liberty to authenticate domain users. WebSphere Application Server Liberty requires that the Microsoft Netlogon service is enabled and running to authenticate these users.

To verify that the Netlogon service is running on the domain controller computer and the computer that is a member of a domain, complete the following steps:

- Right-click Computer and select Manage.
- In the navigation tree view, click Server Manager > Configuration > Services.
- Verify that the Netlogon service is started.
If the service has a Stopped/Disabled status on the domain controller computer, you must also restart the computer that is a member of a Windows domain after you start the service on the domain controller computer.

Installing Db2 by using a Windows domain user account

Before you install IBM Spectrum® Control in a Windows domain, you must install Db2® and register the Db2 license key.

The following table shows the privileges that are required for users associated with installing and running Db2 in a Windows domain environment:

Table 1. User privileges related to Db2 in a Windows domain environment

User role	Required privileges
Installing Db2 on the target Windows domain member machine such that the Db2 instance owner (typically db2admin) is a domain user and the Db2 groups (typically DB2ADMNS and DB2USERS) are domain groups	Local Administrator and Domain Administrator The Db2 installation needs to access and update the Db2 user and group information in the Active Directory that is running on the Windows Domain Controller machine.
Db2 instance owner (typically db2admin) that runs Db2 after Db2 is installed	Local Administrator

Tip: It is recommended that you manually create the Db2 user (the user who will own the Db2 instance that gets created when Db2 is installed, typically *db2admin*), and the Db2 groups (typically *DB2ADMNS* and *DB2USERS*) in the Active Directory on the Domain Controller before you install Db2 on the Windows domain member machine. The Db2 user *must* be a member of both Db2 groups in the Active Directory and a member of the local Administrators group on the Windows domain member machine. To install Db2 by using a Windows domain user account, complete the following steps:

- Log on to Windows as a user that has the correct privileges.
- In Windows Explorer, go to the directory where Db2 installation image is located.
- Right-click **setup.exe** and select Run as administrator.
- Complete the Db2 Setup program.
- On the Set user information for the Db2 Administration Server page, in the User Information panel, select your domain.
- Enter the name of the Db2 user (the user who will own the Db2 instance, typically *db2admin*) that you manually created in the Active Directory on the Domain Controller then enter and confirm the corresponding password.
- On the Enable operating system security for Db2 objects page, complete the following steps:
 - Verify that Enable operating system security check box is selected.
 - In the Db2 administrators group and Db2 users group sections, select your domain and then enter the names of the Db2 groups that you manually created in the Active Directory on the Domain Controller.
 - Click **OK** in the dialog warning you that these groups already exist.

Creating a Windows domain common user account for IBM Spectrum Control

You must create a Windows domain common user account before you can install IBM Spectrum® Control in a Windows domain.

Important: You set the common user for IBM Spectrum Control during the initial installation of the product. When that common user is set, it cannot be changed during the life of that IBM Spectrum Control system. For example, if you set the common user to be *db2admin* during the initial installation of IBM Spectrum Control on a Windows operating system, then the *db2admin* user is the constant common user for that particular IBM Spectrum Control system.

To create a domain common user account for IBM Spectrum Control, complete the following steps:

- Log on to the domain controller computer by using a domain administrator user account.
- Click Start > Administrative Tools > Active Directory Users and Computers.
- Right-click on the Users folder and select New > User.
- On the domain controller computer, click Start > All Programs > Administrative Tools > Active Directory Users and Computers.
- Right-click on the domain user you created and click Properties > Member Of.
- On the Member of tab, click Add and add DB2ADMNS.
- Click Check Names.
- If the name you entered is located, click OK.
- Repeat steps 4 - 8 and add the DB2USERS group.
Tip: The Domain Users group exists.
- Follow the prompts to create the domain common user account for IBM Spectrum Control.
- Log on to the domain client computer on which you want to install IBM Spectrum Control by using a domain administrator user account.
When you log in to computer that is a member of a Windows domain by using the domain administrator user account, and you encounter the 'DB2® UDB SQL5005C System Error' message, complete the following steps:
 - Log in to the computer that is a member of a domain by using the local user name.
 - Click Server Manager > Configuration > Local Users and Groups > Groups > DB2ADMNS > Properties > Add.
 - Add *domain\domain user name*.

- d. Click Check Names.
- e. If the user name that you checked for is underlined, click Apply.
12. Open Server Manager.
13. In the navigation tree, click Server Manager, Configuration, Local Users and Groups, Groups.
14. Add the domain common user name that you created in step 3 (**commonid**) to the local administrator group.

Restriction: When you install IBM Spectrum Control by using a Windows domain user account, the *sAMAccountName* and *userPrincipalName* (the User login name) Active Directory attributes in the Windows domain user account cannot contain spaces. The *sAMAccountName* must have the same value as the part of the *userPrincipalName* Active Directory attribute that precedes the @ character.

Granting Db2 SYSADM authority to a Windows domain user account

If a Windows domain user account is used to install IBM Spectrum® Control, the user account may not have the Db2® SYSADM authority, because Db2 goes to the domain controller computer to list the groups. Before you install IBM Spectrum Control, you must grant the Windows domain user accounts the Db2 SYSADM authority.

To enter the commands, and Db2 SYSADM authority, complete the following steps:

1. Click Start, Command Prompt.
2. Right-click Command Prompt and select Run as administrator.
3. Enter these commands:

```
db2cmd
db2set -g DB2_GRP_LOOKUP=local,TOKENLOCAL
db2 force application all
db2stop
db2start
db2set -all
```

After you run the commands, you can log in to IBM Spectrum Control by using a user account that has Db2 SYSADM authority.

If you are logged in by using your Windows domain user account, you may not have the authorization to run the **db2 force application all** command.

Important: When you grant Db2 SYSADM authority, and the **db2 force application all** command does not work, you can enter the **db2stop** command instead of the **db2 force application all** command.
 For more information about acquiring user group information in the Windows operating system, see [Using an access token to acquire users' group information \(Windows\)](#).
 For more information about the sysadmin_group parameter, see [System administration authority \(SYSADM\)](#).

Related reference

- [Installing IBM Spectrum Control on a Windows domain](#)

Required user privileges in installation scenarios

When you are planning to install IBM Spectrum® Control, there are users to consider for the installation.

These are the users to consider:

- Application administrator
 - The user who logs in to the operating system and does the installation of IBM Spectrum Control.
- Common user
 - The user that you set during the installation of IBM Spectrum Control. The user name that is used to run IBM Spectrum Control.
 - Important: You set the common user for IBM Spectrum Control during the initial installation of the product. When that common user is set, it cannot be changed during the life of that IBM Spectrum Control system. For example, if you set the common user to be *db2admin* during the initial installation of IBM Spectrum Control on a Windows operating system, then the *db2admin* user is the constant common user for that particular IBM Spectrum Control system.
- Db2® user
 - The user that you set during the installation of IBM Spectrum Control. The user name that is used to install and operate the database repository for IBM Spectrum Control.

There are user privileges that are required in various IBM Spectrum Control installation scenarios.

Important: The required user privileges for Common User and Db2 user *cannot* be changed or removed after you complete the installation of IBM Spectrum Control. You can change the required user privileges for the Application administrator after that user completes the installation of IBM Spectrum Control, but only if the Application administrator is *not* set as the Common User or Db2 user in IBM Spectrum Control.
 The following table shows the required user privileges in IBM Spectrum Control installation scenarios:

Table 1. Required user privileges for an IBM Spectrum Control installation

Installation scenario	Required privilege
-----------------------	--------------------

Installation scenario	Required privilege
Setting <i>only</i> the Common User, single-server installation, <i>not</i> in a Windows domain environment	<p>Application administrator:</p> <p>Windows operating system: You must be an Administrator and must have the <i>Debug programs</i> privilege and must <i>not</i> have the <i>Deny access to this computer from the network</i> privilege in the Windows operating system security policy. These privileges can be found in Administrative Tools -> Local Security Policy -> Local Policies -> User Rights Assignment</p> <p>Linux® operating system: You must be a root user or a non-root user with sudo privileges.</p> <p>AIX® operating system: You must be a root user or a non-root user with sudo privileges.</p> <p>Common user:</p> <p>Windows operating system: A member of the local Administrators group, a member of the local Db2 groups (typically DB2ADMNS and DB2USERS).</p> <p>Linux operating system: A member of the root group, a member of the Db2 group (typically db2iadm1).</p> <p>AIX® operating system: member of the system group, member of the Db2 group (typically db2iadm1).</p>
Setting the Common User and Db2 user, single-server installation, <i>not</i> a Windows domain environment	<p>Application administrator:</p> <p>Windows operating system: You must be an Administrator and must have the <i>Debug programs</i> privilege and must <i>not</i> have the <i>Deny access to this computer from the network</i> privilege in the Windows operating system security policy. These privileges can be found in Administrative Tools -> Local Security Policy -> Local Policies -> User Rights Assignment</p> <p>Linux operating system: You must be a root user or a non-root user with sudo privileges.</p> <p>AIX operating system: You must be a root user or a non-root user with sudo privileges.</p> <p>Common User:</p> <p>Windows operating system: A member of the local Administrators group.</p> <p>Linux operating system: A member of the root group.</p> <p>AIX operating system: A member of the system group.</p> <p>Db2 user:</p> <p>Windows operating system: A member of the local Db2 groups (typically DB2ADMNS and DB2USERS).</p> <p>Linux operating system: A member of the Db2 group (typically db2iadm1).</p> <p>AIX operating system: A member of the Db2 group (typically db2iadm1).</p>
Setting <i>only</i> the Common User, single-server installation, in a Windows domain environment with Db2 installed so the Db2 instance owner (typically db2admin) is a domain user and the Db2 groups (typically DB2ADMNS and DB2USERS) are domain groups.	<p>Application administrator:</p> <p>You must be an Administrator and must have the <i>Debug programs</i> privilege and must <i>not</i> have the <i>Deny access to this computer from the network</i> privilege in the Windows operating system security policy. These privileges can be found in Administrative Tools -> Local Security Policy -> Local Policies -> User Rights Assignment</p> <p>Common user:</p> <p>A member of the local Administrators group, a member of the domain Db2 groups (typically DB2ADMNS and DB2USERS; these groups are stored in the domain Active Directory).</p>
Setting the Common User and Db2 user, single-server installation, in a Windows domain environment with Db2 installed so the Db2 instance owner (typically db2admin) is a domain user and the Db2 groups (typically DB2ADMNS and DB2USERS) are domain groups.	<p>Application administrator:</p> <p>You must be an Administrator and must have the <i>Debug programs</i> privilege and must <i>not</i> have the <i>Deny access to this computer from the network</i> privilege in the Windows operating system security policy. These privileges can be found in Administrative Tools -> Local Security Policy -> Local Policies -> User Rights Assignment</p> <p>Common User: A member of the local Administrators group.</p> <p>Db2 user: A member of the domain Db2 groups (typically DB2ADMNS and DB2USERS; these groups are stored in the domain Active Directory).</p>

Installation scenario	Required privilege
Setting <i>only</i> the Db2 user, multiple server installation (Database repository), <i>not</i> in a Windows domain environment.	<p>Application administrator:</p> <p>Windows operating system:</p> <p>You must be an Administrator and must have the <i>Debug programs</i> privilege and must <i>not</i> have the <i>Deny access to this computer from the network</i> privilege in the Windows operating system security policy. These privileges can be found in Administrative Tools -> Local Security Policy -> Local Policies -> User Rights Assignment</p> <p>Linux operating system: Must be a root user or a non-root user with sudo privileges.</p> <p>AIX operating system: Must be a root user or a non-root user with sudo privileges</p> <p>Db2 user:</p> <p>Windows operating system: A member of the local DB2® groups (typically DB2ADMNS and DB2USERS).</p> <p>Linux operating system: A member of the DB2 group (typically db2iadm1).</p> <p>AIX operating system: A member of the DB2 group (typically db2iadm1).</p>
Setting <i>only</i> the Common User, multiple server installation (IBM Spectrum Control Servers), <i>not</i> in a Windows domain environment.	<p>Application administrator:</p> <p>Windows operating system:</p> <p>You must be an Administrator and must have the <i>Debug programs</i> privilege and must <i>not</i> have the <i>Deny access to this computer from the network</i> privilege in the Windows operating system security policy. These privileges can be found in Administrative Tools -> Local Security Policy -> Local Policies -> User Rights Assignment</p> <p>Linux operating system: Must be a root user or a non-root user with sudo privileges.</p> <p>AIX operating system: Must be a root user or a non-root user with sudo privileges.</p> <p>Common user:</p> <p>Windows operating system: A member of the local Administrators group.</p> <p>Linux operating system: A member of the root group.</p> <p>AIX operating system: A member of the system group.</p>
Setting <i>only</i> the Db2 user, multiple server installation (Database repository), in a Windows domain environment with Db2 installed so the Db2 instance owner (typically db2admin) is a domain user and the Db2 groups (typically DB2ADMNS and DB2USERS) are domain groups.	<p>Application administrator:</p> <p>You must be an Administrator and must have the <i>Debug programs</i> privilege and must <i>not</i> have the <i>Deny access to this computer from the network</i> privilege in the Windows operating system security policy. These privileges can be found in Administrative Tools -> Local Security Policy -> Local Policies -> User Rights Assignment</p> <p>Db2 user:</p> <p>A member of the domain Db2 groups (typically DB2ADMNS and DB2USERS; these groups are stored in the domain Active Directory).</p>
Setting <i>only</i> the Common User, multiple server installation (IBM Spectrum Control Servers), in a Windows domain environment.	<p>Application administrator:</p> <p>You must be an Administrator and must have the <i>Debug programs</i> privilege and must <i>not</i> have the <i>Deny access to this computer from the network</i> privilege in the Windows operating system security policy. These privileges can be found in Administrative Tools -> Local Security Policy -> Local Policies -> User Rights Assignment</p> <p>Common user:</p> <p>A member of the local Administrators group.</p>

Planning for configuration

Use this information to plan your IBM Spectrum® Control environment. The information includes general configuration guidelines, TCP/IP ports, user names and user rights, and SMI-S support.

- [Planning for capacity](#)
The following guidelines can help you determine the capacity for IBM Spectrum Control.
- [Planning for IBM Spectrum Control authentication and authorization](#)
An operating system user name is required to install and log on to IBM Spectrum Control for the first time. After you install IBM Spectrum Control, you can assign roles to users. Roles determine the product functions that are available to users.
- [User names and passwords](#)
Several user names and passwords are required to install, configure, and use IBM Spectrum Control. There are also some requirements and limitations that you must understand before you install IBM Spectrum Control.
- [Ports used by IBM Spectrum Control](#)
When you install IBM Spectrum Control, the ports must be opened through the firewall. You must disable the firewall program or open the ports to allow incoming requests to the IBM Spectrum Control ports. Review these ports before you install IBM Spectrum Control.
- [Planning for multipath subsystem device drivers](#)
The subsystem device driver (SDD) is a software solution for multiple configuration environments in supported storage resources.

- [Planning for Storage Resource agents](#)
Use Storage Resource agents to collect asset and configuration information about servers.
- [Planning for Internet Protocol Version 6](#)
IBM Spectrum Control supports Internet Protocol Version 6 (IPv6) for communication between its components. The key IPv6 enhancement is the expansion of IP address spaces from 32 bits (up to 15 characters in length) to 128 bits (up to 45 characters in length).
- [Planning to use LDAP for IBM Spectrum Control authentication](#)
The Lightweight Directory Access Protocol (LDAP) is an application protocol that you can use to query and modify directory services running over TCP/IP. The IBM Spectrum Control installation program establishes a default authentication configuration using the federated repositories feature of the IBM® WebSphere® Application Server Liberty. You can configure IBM Spectrum Control for LDAP authentication as a post-installation activity.

Planning for capacity

The following guidelines can help you determine the capacity for IBM Spectrum® Control.

About this task

Consider the following guidelines when you determine your capacity requirements:

Database repository

For best optimal read/write performance, install IBM Spectrum Control, Db2® data, and the Db2 logs on separate physical disks, which are separate from the host server's operating system. Of course, how you configure your environment is based on your storage technology.

IBM Spectrum Control server

The IBM Spectrum Control server must meet the following requirements:

- The IBM Spectrum Control server must have dual 3.2 GHz processors with a minimum of 8 GB of random access memory (RAM).
- The IBM Spectrum Control server must be a dedicated computer for IBM Spectrum Control operations and not shared with other applications.
- Paging space: On AIX® operating systems, the paging space must be at least half or, if possible, the same amount as the physical memory. You must periodically monitor paging space usage by running the `lsps -s` command and increasing the space if the percentage used is higher than 50%.

LSI SMI-S Provider

IBM Spectrum Control supports up to five storage systems with the LSI SMI-S Provider 1.3 or later.

Planning for IBM Spectrum Control authentication and authorization

An operating system user name is required to install and log on to IBM Spectrum® Control for the first time. After you install IBM Spectrum Control, you can assign roles to users. Roles determine the product functions that are available to users.

User requirements for installing IBM Spectrum Control

The user name that you use when you run the IBM Spectrum Control installation program must belong to the following operating system groups. The user that you define as the common user during the IBM Spectrum Control installation must also belong to these groups. The IBM Spectrum Control installation program automatically maps these groups to the IBM Spectrum Control Administrator role.

Table 1. Operating system groups that are mapped to the Administrator role

Operating System	Group
Windows	Administrators
AIX®	system
Linux®	root

The user name that is used to install the database repository must belong to one of the following Db2® groups. The user name that is used to install the database repository depends on your installation configuration. For example, if you installed IBM Spectrum Control on a single server and you did not enter separate user information for Db2, the database repository is installed by using the common user name. Therefore, the common user must have these Db2 privileges.

Table 2. Db2 groups that are required to install the database repository

Operating System	Group
Windows	DB2ADMNS
AIX and Linux	db2iadm1

The IBM Spectrum Control installation program establishes a default authentication configuration by using the federated repositories feature of the WebSphere® Application Server Liberty.

If you want to change the user authentication configuration, see [Configuring user authentication](#).

For more information about installing IBM Spectrum Control on a Windows domain, see [Planning to install IBM Spectrum Control in a Windows domain](#) and [Installing IBM Spectrum Control on a Windows domain](#).

Defining roles for IBM Spectrum Control

After you install IBM Spectrum Control, you can assign roles for IBM Spectrum Control users. These roles are predefined in IBM Spectrum Control and are assigned at the group level. The role determines the authorization level for the users who are in the group. The authorization level determines the functions in IBM Spectrum Control that the user can access.

For example, if a user is in a group that is assigned to the Administrator role, the user has full access to all IBM Spectrum Control functions. However, if a user is in a group that is assigned to the Monitor role, the user has access to only a limited set of functions.

For more information about roles and groups, see [Role-based authorization](#).

User names and passwords

Several user names and passwords are required to install, configure, and use IBM Spectrum® Control. There are also some requirements and limitations that you must understand before you install IBM Spectrum Control.

A worksheet is also provided to help you plan and document the user names and passwords that you use with IBM Spectrum Control. Keep the completed worksheet in a safe, secure place. Before you install IBM Spectrum Control, you must understand your local security policies and requirements. These policies can impose certain standards that you must consider when you create user names and passwords to comply with requirements.

- [User name and password requirements](#)
This information helps you create valid user names and passwords for IBM Spectrum Control and understand any special character restrictions.
- [Db2 user names and passwords](#)
These are the rules for using Db2® user names and passwords.
- [Required user roles for monitoring resources](#)
When you add storage systems, switches, and hypervisors for monitoring, you must provide a user name and password for logging in to those resources. The role or user group that is assigned to the user name determines the data collection and storage functions that you can use.
- [Worksheet for user names and passwords](#)
Use this worksheet to document the user names and passwords that you create when you install and administer IBM Spectrum Control. Space is provided for you to list more user names and passwords that you create when you configure more devices in IBM Spectrum Control.

Related reference

- [User name and password requirements](#)
- [Db2 user names and passwords](#)
- [Worksheet for user names and passwords](#)

User name and password requirements

This information helps you create valid user names and passwords for IBM Spectrum® Control and understand any special character restrictions.

[Table 1](#) lists the user names and passwords that are used with IBM Spectrum Control and specifies the following criteria:

- User names must be local or a Windows domain account or an LDAP account
Restriction: If you use LDAP, IBM Spectrum Control must be configured for LDAP authentication.
- Valid user name characters must comply with the minimum and maximum length requirements
- Valid password characters must comply with the minimum and maximum length requirements

Restriction: If the IBM Spectrum Control common user name is the same as the IBM Spectrum Control Db2® user name (for example, if you did not use a custom Db2 user name to install IBM Spectrum Control), the corresponding password must be valid as a Db2 password and as a WebSphere® Liberty password. For more information about valid Db2 and WebSphere Liberty passwords, see [Table 1](#). This restriction applies during a IBM Spectrum Control installation and when you use the IBM Spectrum Control change password tool.

For more information about setting passwords on the Windows operating system, see the see the Microsoft Windows Dev Center - Desktop website at <http://msdn.microsoft.com/en-us/library/windows/desktop/ms675092%28v=vs.85%29.aspx>. Search for the *SAM-Account-Name* attribute.

Table 1. Valid characters for user names and passwords

User name and password for	Windows domain OK?	LDAP OK?	Valid characters
IBM Spectrum Control installation	Yes	No	Not applicable
IBM Spectrum Control common user	Yes Restriction: You must follow these naming conventions: <ul style="list-style-type: none">• Windows domain user name: <i>domain_name\user_name format</i>• Local user name: <i>machine_name\username format</i>	No	<ul style="list-style-type: none">• A through Z (uppercase characters)• a through z (lowercase characters)• 0 through 9 (numeric characters)• Special characters: '_' and '~' <p>Restriction: The user name cannot contain spaces, \$ and ! signs, must have at least 1 character and must not start with a number or underscore. On Windows operating systems, the only valid special characters are ` , ~ , # , % , (,) , - , _ , { , } , ' , and .</p>
IBM Spectrum Control common user passwords	Yes	Yes	<ul style="list-style-type: none">• A through Z (uppercase characters)• a through z (lowercase characters)• 0 through 9 (numeric characters)• Special characters: [,] , ? , ` , ~ , ! , (,) , - , _ , * , % , + , and . <p>Restriction: Passwords cannot contain spaces, must have at least 1 character and must not start with a number, dash character {-}, or underscore.</p>

User name and password for	Windows domain OK?	LDAP OK?	Valid characters
Db2 administrator, Db2 user	Yes Restriction: You must follow these naming conventions: <ul style="list-style-type: none"> Windows domain user name: <code>domain_name\user_name format</code> Local user name: <code>machine_name\user_name format</code> 	Yes	<ul style="list-style-type: none"> A through Z (uppercase characters) a through z (lowercase characters) 0 through 9 (numeric characters) Special characters: ` , @ , # , % , ^ , - , () , _ , { } " . <p>Restriction: User names and passwords cannot contain spaces and must have at least 1 character. User names also must not start with a number or underscore</p>
Db2 passwords	Yes	Yes	<ul style="list-style-type: none"> A through Z (uppercase characters) a through z (lowercase characters) 0 through 9 (numeric characters) Special characters: \$, ~ , @ , # , (,) , - , _ , { , } , [,] , ? , * , % , + , and . <p>Restriction: Passwords cannot contain spaces, must have at least 1 character and must not start with a number or underscore.</p>
NAS Filer user name	No	No	<ul style="list-style-type: none"> A through Z (uppercase characters) a through z (lowercase characters) 0 through 9 (numeric characters) Special characters: ` ~ # % ^ & () - _ { } ' . <p>Restriction: User names and passwords cannot contain spaces and must have at least 1 character.</p>
NAS Filer password	No	No	<ul style="list-style-type: none"> A through Z (uppercase characters) a through z (lowercase characters) 0 through 9 (numeric characters) Special characters: ` ~ @ # % ^ & * () - _ = + [] { } \ ; : ' " , . < > / ? <p>Restriction: User names and passwords cannot contain spaces and must have at least 1 character.</p>

Password and name restrictions for monitored resources:

- Cisco switches: The password that IBM Spectrum Control uses to connect to a Cisco switch can't contain these special characters: < > .
- XIV® storage systems: The password that IBM Spectrum Control uses to connect to an XIV storage system can't contain this special character: & .
- When you add a resource for monitoring, you must enter a user name and password that is used to connect to that resource. In rare cases, special characters in the user name or password might not be allowed by IBM Spectrum Control. If you can't add a resource for monitoring because of restricted characters, try changing its credentials and adding it again. Then, open a ticket for the character limitation so that it can be addressed in a future update.
- Performance metadata can't be collected about a device that includes a forward slash (/) or backslash (\) in its name. Before you add a device for monitoring, remove any slashes in its name.

Db2 user names and passwords

These are the rules for using Db2® user names and passwords.

Db2 user names and passwords must follow these rules:

- UNIX user names and passwords cannot be more than eight characters long. They cannot begin with a numeric digit or end with \$.
- Group and instance names can contain 1 to 8 characters.
- Names cannot be any of the following:
 - USERS
 - ADMINS
 - GUESTS
 - PUBLIC
 - LOCAL
- Names cannot begin with:
 - IBM®
 - SQL
 - SYS
- Names cannot include accented characters.
- UNIX users, groups, and instance names must be lowercase.

Required user roles for monitoring resources

When you add storage systems, switches, and hypervisors for monitoring, you must provide a user name and password for logging in to those resources. The role or user group that is assigned to the user name determines the data collection and storage functions that you can use.

The following roles are associated with the user names that IBM Spectrum® Control uses to log in to resources. Specify user names when you add a resource for monitoring. These roles are different from IBM Spectrum Control roles, which are assigned to users that log in to IBM Spectrum Control. For more information about IBM

Spectrum Control roles, see [Role-based authorization](#).

Any roles that are not listed, but include the privileges of the roles that are listed, can also be used for monitoring resources.

Table 1. Required roles for storage system, switch, and hypervisor users

Resource	Required role for performance monitors	Required role for probes	Required role for provisioning, optimizing, and transforming
DS8000®	Monitor role or higher	Monitor role or higher	Administrator, Physical operator, or Logical operator
SAN Volume Controller	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For 8.3.1.2 or later, any role but some limitations might apply. Learn more	Monitor role or higher	Administrator
Storwize® V3500 Storwize V3700 Storwize V5000 Storwize V7000	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For 8.3.1.2 or later, any role but some limitations might apply. Learn more	Monitor role or higher	Administrator
Storwize V7000 Unified (block storage)	Administrator	Monitor role or higher	Administrator
FlashSystem 5000 FlashSystem 5100 FlashSystem 7200 FlashSystem 9100 FlashSystem 9200 FlashSystem V9000	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For 8.3.1.2 or later, any role but some limitations might apply. Learn more	Monitor role or higher	Administrator
IBM Spectrum Virtualize for Public Cloud	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For 8.3.1.2 or later, any role but some limitations might apply. Learn more	Monitor role or higher	Administrator
Storwize V7000 Unified (file storage)	Performance monitoring is not available	Monitor	Administrator, or Export Administrator + Storage Administrator.
XIV®	Monitor role or higher	Any role	Storage administrator
IBM Spectrum Accelerate IBM FlashSystem® A9000 IBM FlashSystem A9000R	Monitor role or higher	Any role	Not available
IBM FlashSystem 900	Monitor role or higher	Monitor role or higher	Not available
IBM Spectrum Scale (file storage)	Any role	Root or non-root*	Not available
IBM Spectrum Scale (object storage)	Performance monitoring is not available	The Keystone admin role. Information is collected only for the object storage accounts and containers that the user has access to. If you want to monitor all accounts and containers, the user must also be assigned the role that is defined in the reseller_admin_role configuration option in the Swift proxy server. The default value for the reseller_admin_role option is ResellerAdmin.	Not available

Resource	Required role for performance monitors	Required role for probes	Required role for provisioning, optimizing, and transforming
IBM® Cloud Object Storage	Performance monitoring is not available	Operator, System Administrator, or Super User role	Not available
Dell EMC storage systems	Operator role or higher	Operator role or higher	Not available
Hitachi storage systems	In Hitachi Device Manager: The user for the Export Tool must have the Storage Administrator (Performance Management) role	In Hitachi Command Suite: ViewGroup or higher and AdminGroup permission.	Provisioning is not available
NetApp storage systems	Operator role or higher	Operator role or higher	Not available
Pure storage systems	Role with read-only permission or higher	Role with read-only permission or higher	Not available
Cisco switches	Network-admin	Network-admin	Not applicable
Brocade switches with Fabric OS 8.2.1 or later	User or admin role that has the chassis-role permission	User or admin role that has the chassis-role permission	Not applicable
Brocade switches with a Fabric OS version earlier than 8.2.1	Administrator user in Brocade Network Advisor (BNA)	Administrator user in BNA	Not applicable
Hypervisors (such as ESX, ESXi, and vCenter Servers)	Performance monitoring is not available	A role that has permission to browse through data stores. For example: Administrator role or Virtual Machine Power® User role. For more information about roles and permission to browse data stores, see Checking permissions to browse data stores .	Any role
Rollup servers	Performance monitoring is not available	Administrator	Not available
Note: * You can add IBM Spectrum Scale and GSS systems as a non-root user, but that user must have privileges to run a set of specified administration commands using the sudo command on the cluster node. For more information, see Monitoring IBM Spectrum Scale without requiring root privileges .			

Worksheet for user names and passwords

Use this worksheet to document the user names and passwords that you create when you install and administer IBM Spectrum® Control. Space is provided for you to list more user names and passwords that you create when you configure more devices in IBM Spectrum Control.

Table 1. User names and passwords for a IBM Spectrum Control installation

Item	Description	Your input
IBM Spectrum Control and Db2® installation user <ul style="list-style-type: none"> Windows: local administrator UNIX: root 		
Db2 administrator user name and password	The Db2 administrator user name and password are created when you install Db2 and is required to install IBM Spectrum Control. On the Windows operating system, this user name must be a member of the DB2ADMNS group and Administrators group. On UNIX, the user name must be the instance owner of the instance you want to use. This user name and password are created when you install Db2.	
LDAP bind user name	This user name has the authority to log in and connect to the LDAP server for authentication. It is not needed if the LDAP server allows anonymous bind. This user name and password must be provided to you by your LDAP administrator.	
Storage subsystem user name for access from IBM Spectrum Control	This user name is the user name that IBM Spectrum Control uses to log in to a storage device.	
CIMOM or SMI provider login name	This user name has log in access to the CIMOM or SMI provider software that interfaces with a storage device or a switch in a fabric.	
NAS filer user name and password	IBM Spectrum Control uses this user name to log in to a NAS device.	

Related reference

- [User name and password requirements](#)

Ports used by IBM Spectrum Control

When you install IBM Spectrum® Control, the ports must be opened through the firewall. You must disable the firewall program or open the ports to allow incoming requests to the IBM Spectrum Control ports. Review these ports before you install IBM Spectrum Control.

Restriction: IBM Spectrum Control uses TLS 1.2 protocol for communicating on ports. It does not use TLS 1.3, and TLS 1.1 and 1.0 are disabled by default for increased security.

For information about how to re-enable TLS 1.1 and 1.0 for IBM Spectrum Control ports, see [Enabling TLS 1.1 and 1.0 for IBM Spectrum Control ports](#).

IBM Spectrum Control ports used for TCP/IP

Table 1 lists the IBM Spectrum Control default ports.

All ports in Table 1 are configured during installation using the installation program. For a silent mode installation, use the varTPCPortRangeSP parameter in the installation response file.

Table 1. IBM Spectrum Control ports used by TCP/IP for incoming communication

Ports	Component and port information
9549	Data server 9549 listens for communications from the Device server, CLI, and Storage Resource agents.
9550-9551 9553 9572	Device server 9550-9551 listens for communications from the Data server, Alert server, and Web server. 9553 listens for logging commands. 9572 listens for IBM® WebSphere® Application Server Liberty administrative commands.
9556 9570 9571	Alert server 9556 listens for communications from the Device server. 9570 listens for logging commands. 9571 listens for WebSphere Application Server Liberty administrative commands.
9510 9567	Storage Resource agent 9510 listens for communications from the Data server if the Storage Resource agent is deployed in daemon mode. The only exception is the Storage Resource agent that is deployed on the IBM Spectrum Control server at installation time, which listens on port 9567 instead. If the Storage Resource agent is deployed in non-daemon mode, the Storage Resource agent does not need any listening ports. The Data server uses RXA to communicate with the Storage Resource agent in non-daemon mode. RXA uses the following default ports for the protocols: 22 SSH 512 Remote Execution (REXEC) 514 Remote Shell (rsh) 445 Windows SMB protocol The Storage Resource agent sends responses back to the Data server on port 9549, in both daemon mode and non-daemon mode.
9554 9568 9569	Web server 9554 is used by WebSphere Application Server Liberty. The following ports are used for the GUI: <ul style="list-style-type: none">• 9568 (non-secure)• 9569 (secure) 9569 is used for the REST API.
9562	Export server 9562 listens for communications from the Web server.
25000 50000	Db2® By default, Db2 listens on port 25000 for communication from IBM Spectrum Control and other clients. Port 50000 can also be used. The default port can be changed only during the Db2 installation.

Restrictions: The following restrictions apply to the ports that are listed in [Table 1](#):

- The ports apply only to a fresh installation of IBM Spectrum Control. If you upgrade from an earlier version of IBM Spectrum Control that used different ports, those ports are reused for your installation.
- The ports are used for incoming communication and are configured during installation. *Do not* change these ports after your installation is complete.

Other IBM Spectrum Control IP ports

Table 2. List of IP ports that are used by IBM Spectrum Control

Resource and ports	Port information
<p>Resources that run IBM Spectrum Virtualize</p> <p>Ports: 22, 5989 (both unidirectional)</p>	<p>For native interface:</p> <ul style="list-style-type: none"> • 22 listens for SSH communication from IBM Spectrum Control. • 5989 listens for other communication from IBM Spectrum Control. You can upload the SSH key once at setup time using the 5989 port. <p>By default, the IBM Spectrum Control Device server creates a connection between a IBM Spectrum Virtualize resource and a local port to listen for events from the IBM Spectrum Virtualize resource. The default range is 49152 to 65535. These ports are on the Device server only; you do not need to configure the firewall for them. The ports are used for a SSH tunnel to connect to the event port on the IBM Spectrum Virtualize resource. All communication to the IBM Spectrum Virtualize resource runs through the SSH connection. No other ports need to be open for IBM Spectrum Virtualize events.</p> <p>To change to a custom port range, use the CLI setdscfg command to change the values of the following parameters:</p> <ul style="list-style-type: none"> • NAPI.SVCEventListenerLPortBase • NAPI.SVCEventListenerLPortMax <p>After resetting the value, restart the Device server.</p> <p>For example, to set the range from 50100 to 65000, run these commands from the <i>installation_dir/cli/</i> subdirectory:</p> <p>For Windows operating system:</p> <pre>tpctool.bat setdscfg -user <user name> -pwd <password> -property NAPI.SVCEventListenerLPortBase 50100 tpctool.bat setdscfg -user <user name> -pwd <password> -property NAPI.SVCEventListenerLPortMax 65000</pre> <p>For AIX® and Linux® operating systems:</p> <pre>./tpctool.sh setdscfg -user <user name> -pwd <password> -property NAPI.SVCEventListenerLPortBase 50100 ./tpctool.sh setdscfg -user <user name> -pwd <password> -property NAPI.SVCEventListenerLPortMax 65000</pre>
<p>FlashSystem 900</p> <p>Ports: 22, 5989, 161 (UDP) (all unidirectional)</p>	<p>For native interface:</p> <ul style="list-style-type: none"> • 22 listens for SSH communication from IBM Spectrum Control. • 5989 listens for other communication from IBM Spectrum Control. You can upload the SSH key once at setup time using the 5989 port. <p>For SNMP:</p> <ul style="list-style-type: none"> • 161 (UDP) listens for SNMP communications from IBM Spectrum Control for performance monitoring. • The SNMP agent must be enabled for a storage system before IBM Spectrum Control can collect its performance data.
<p>DS8000®</p> <p>Ports: 1751, 1750, 1755 (all bidirectional)</p> <p>Ports: 8451, 8452 (both unidirectional)</p>	<p>Native interface:</p> <p>1751 (default) listens for communication from IBM Spectrum Control and other clients. If 1751 is unavailable, uses 1750.</p> <p>1751 (default) sends hardware management console (HMC) information and events to IBM Spectrum Control. If 1751 is unavailable, uses 1750.</p> <p>1755 sends and receives data for HTC when logs are being offloaded.</p> <p>8451, 8452 DS8000 listens on these ports for communication with the DS8000 GUI.</p>
<ul style="list-style-type: none"> • IBM XIV® Storage System • IBM Spectrum Accelerate • FlashSystem A9000 • FlashSystem A9000R <p>Port: 7778 (bidirectional)</p>	<p>The storage systems listen for communication from IBM Spectrum Control and other clients, and also send communications to IBM Spectrum Control.</p>

Resource and ports	Port information
IBM Spectrum Scale <ul style="list-style-type: none"> File storage ports: 22, 9084 (both unidirectional) Object storage ports: 5000, 8080, 35357 (all unidirectional) 	<ul style="list-style-type: none"> File storage: <ul style="list-style-type: none"> 22 listens for SSH communication from IBM Spectrum Control. 9084 listens for performance data collection on the IBM Spectrum Scale cluster node where the collector component is running. To collect performance metadata, complete the following actions: <p>IBM Spectrum Scale 5.1.0 and earlier:</p> <p>Ensure that the queryinterface property is set to "0.0.0.0" in the ZIMonCollector.cfg file on the cluster node.</p> <pre>queryinterface="0.0.0.0"</pre> <p>IBM Spectrum Scale 5.1.1 and later:</p> <p>Ensure that the zimmon/ZIMonCollector.cfg file on the cluster node includes the following properties:</p> <pre>fallbackqueryinterface = "0.0.0.0" # "0.0.0.0" to allow remote connections (or ":::0" for IPv6) fallbackqueryport = "9084"</pre> <p>For more information, see Configuring the collection of performance data for IBM Spectrum Scale.</p> Object storage: <ul style="list-style-type: none"> 5000, 8080, 35357 listen for data about any object storage or OpenStack Swift configurations on the IBM Spectrum Scale cluster node where the collector component is running.
IBM Cloud Object Storage Port: 443 (unidirectional)	443 listens for HTTPS communication from IBM Spectrum Control.
Dell EMC Unity Port: 443 (unidirectional)	443 listens for communication from IBM Spectrum Control.
Other Dell EMC storage systems Port: 5991 (unidirectional)	<p>5991 listens for events that are coming from Dell EMC CIMOMs .</p> <p>To change this port, use the CLI setdscfg command and change the value for the Event.ListenerPort parameter. Then, restart the Device server.</p> <p>For example, to set the port to 7200, run this command from the <i>installation_dir/cli/</i> subdirectory:</p> <p>For Windows operating system:</p> <pre>tpctool.bat setdscfg -user <user name> -pwd <password> -property Event.ListenerPort 7200</pre> <p>For AIX and Linux operating systems:</p> <pre>./tpctool.sh setdscfg -user <user name> -pwd <password> -property Event.ListenerPort 7200</pre> <p>You can also modify the protocol for listening to events that are coming from Dell EMC CIMOMs. By default, this protocol is set http.</p> <p>To change this protocol, use the CLI setdscfg command and change the value for the Event.ListenerProtocol parameter. Then, restart the External Process component.</p> <p>For example, to set the protocol to https, run this command from the <i>installation_dir/cli/</i> subdirectory:</p> <p>For Windows operating system:</p> <pre>tpctool.bat setdscfg -user <user name> -pwd <password> -property Event.ListenerProtocol https</pre> <p>For AIX and Linux operating systems:</p> <pre>./tpctool.sh setdscfg -user <user name> -pwd <password> -property Event.ListenerProtocol https</pre>
Hitachi VSP storage systems: <ul style="list-style-type: none"> Hitachi Command Suite: Port: 2443 (unidirectional) Hitachi Device Manager: Port: 2443 (unidirectional) 	2443 listens for communication from IBM Spectrum Control.

Resource and ports	Port information
NetApp ONTAP 9 Port: 22 (unidirectional)	22 listens for communication from IBM Spectrum Control.
Other NetApp storage systems Port: 80 (unidirectional)	80 listens for communication from IBM Spectrum Control using the NetApp Data ONTAP API (NAPI), if the storage system is a file storage system. If the NetApp storage system is a block storage system, the CIM agent ports are used to listen for communication from the Device server.
Pure storage systems Port: 443 (unidirectional)	443 listens for communication from IBM Spectrum Control.
SMI-S providers and CIM agents for switches and managed storage systems Ports: 5988, 5989, 5990 (unidirectional)	5988 listens for non-secure communication from the Device server. 5989 listens for secure communication from the Device server. 5990 IBM Spectrum Control listens on this port for events coming from CIM agents for managed storage systems. To change the port that IBM Spectrum Control uses to listen for events coming from CIM agents for managed storage systems, use the CLI setdscfg command. Change the value for the Indication.PermanentListenerPort parameter then restart the Device server. For example, to set the port to 7100, run this command from the <i>installation_dir/cli/</i> subdirectory: For Windows operating system: <pre>tpctool.bat setdscfg -user <user name> -pwd <password> -property Indication.PermanentListenerPort 7100</pre> For AIX and Linux operating systems: <pre>./tpctool.sh setdscfg -user <user name> -pwd <password> -property Indication.PermanentListenerPort 7100</pre>
Brocade switches, Fabric OS 8.2.1 and later Ports: 80, 443 (unidirectional)	80 listens for HTTP communication from IBM Spectrum Control. 443 listens for HTTPS communication from IBM Spectrum Control.
VMware vCenter Servers Port: 443 (unidirectional)	443 listens for communication from IBM Spectrum Control.
IBM Spectrum Control Device server Ports: 162 (UDP), 5960 (UDP), (both unidirectional) Ports: 49152 (UDP) (bidirectional)	162 (UDP) listens for SNMP traps coming from managed storage systems. 5960 (UDP) listens for forwarded SNMP messages. 49152 (UDP) IBM Spectrum Control uses this port to forward SNMP traps. To change these ports, use these steps: 1. Log on to IBM Spectrum Control as a user with administrative privileges. 2. Stop the Device server. 3. Create a backup of the <i>installation_dir/device/conf/user.properties</i> file. 4. Change the parameter that corresponds to the port to an available port number and save the file. 162 SnmpTrapPort 5960 SnmpForwardedTrapListeningPort 49152 SnmpForwardingPort 5. Start the Device server.
All monitored resources 7 (ECHO port) (bidirectional)	Before you can discover resources, ensure that port 7 (ECHO port) is open on the resources, and Internet Control Message Protocol (ICMP) is available.
IBM Cognos® Analytics Port: 9300 (unidirectional)	Listens for communication from web browsers and other clients.

Related information

- [setdscfg](#)
- [Starting and stopping the IBM Spectrum Control servers](#)

Planning for multipath subsystem device drivers

The subsystem device driver (SDD) is a software solution for multiple configuration environments in supported storage resources.

The subsystem device driver is installed on a host system with the native disk-device driver and provides the following functions:

- Enhanced data availability
- Dynamic input/output (I/O) load balancing across multiple paths
- Automatic path failover protection
- You can download licensed machine code at the same time that applications are running

For the most current support for multipath subsystem device drivers, go to [IBM Spectrum Control - Platform Support: Agents, Servers and Browsers - Multipathing section](#).

Note:

1. The AIX® SDD cannot coexist with SDDPCM on the same system.
2. The Linux® SDD driver is no longer available. Starting with Red Hat® Enterprise Linux 5 and SUSE Linux Enterprise Server 10, only the DM_Multipath is available.
3. The SDD driver is no longer supported on HP, starting with HP-UX 11i 3 with Itanium®.

For more information about multipath subsystem device drivers, see <https://www.ibm.com/support/pages/node/651823>.

For more information about how to install, configure, and use the subsystem device drivers, see <https://www.ibm.com/support/pages/node/651689>.

Upgrading subsystem device drivers

SDD drivers cannot coexist on the same host with the SDDPCM, SDDDSM, or DM_Multipath drivers. You must upgrade from the existing SDD drivers to the SDDPCM or SDDDSM driver. Part of the upgrade process is to unconfigure and remove all SDD vpath devices. After the upgrade and configuration of the devices, the device names might differ from the previous names. Each device is detected by IBM Spectrum® Control as a new device. For information about how to upgrade SDD, see <http://www.ibm.com/support/docview.wss?rs=540&context=ST52G7&uid=ssg1S7000303>.

Dell EMC PowerPath Multipathing

Dell EMC PowerPath Multipathing supports a wide range of servers including cluster servers connected to Dell EMC storage systems. It tunes your storage area network and selects alternate paths for your data if necessary. It also integrates multiple path I/O capabilities, automatic load balancing, and path failover functions. For more information about Dell EMC PowerPath Multipathing, see <http://www.dellemc.com>. Search for **PowerPath**.

Planning for Storage Resource agents

Use Storage Resource agents to collect asset and configuration information about servers.

You can deploy Storage Resource agents from the IBM Spectrum® Control GUI. You must have administrative privileges to deploy Storage Resource agents.

To deploy a Storage Resource agent from the GUI, go to **Servers > Servers**. Click **Add Server**, select **Deploy an agent for full server monitoring**, and select a method for adding a server. You can add a single server by manually entering server information or you can add multiple servers by importing configuration information from a file. If you add multiple servers, a time span is calculated during which the agents are deployed. The agents are deployed at regular intervals during the time span to avoid excessive load on the IBM Spectrum Control server.

For more information about Storage Resource agents, see [Deployment guidelines and limitations for Storage Resource agents](#).

- [Protocol support for Storage Resource agents](#)
When installing the Storage Resource agent, IBM Spectrum Control uses specific protocols for connectivity between the server and agent.

Protocol support for Storage Resource agents

When installing the Storage Resource agent, IBM Spectrum® Control uses specific protocols for connectivity between the server and agent.

When installing the Storage Resource agent, IBM Spectrum Control uses the following protocol for connectivity between the server and agent (listed in order):

1. Secure Shell protocol (SSH).
2. Windows server message block protocol (SMB protocol).
3. Remote execution protocol (REXEC).
4. Remote shell protocol (RSH).

At run time, the connectivity that is used between the server and agent depends on the type of service that is running: On-Demand service (non-daemon service) or run as a service (daemon service).

On-Demand service (non-daemon service)

In this case, connectivity between the server and agent is established by using the same protocols as for installation of the agent: SSH, SMB, REXEC, or RSH.

Run as a service (daemon service)

In this case, connectivity between the server and agent is established by using the secured socket connection. The server and agent have their respective certificates and no additional information is needed besides the certificates and the security that is provided by the SSH protocol.

Note: Running the Storage Resource agent in daemon mode requires only one socket connection to be open and thus simplifies any firewall rules that might be in place. However, the daemon mode always consumes resources (although a very small amount) when idle. In non-daemon mode, the Storage Resource agent only consumes resources when actively working for the IBM Spectrum Control server. Because the non-daemon mode uses the RXA protocol, additional firewall ports and security configuration are required.

The information required for these protocols is as follows:

SSH protocol

There are two cases where you can use the SSH protocol:

- You supply the user ID and password to connect to the server by using the SSH protocol. This user ID must have administrative privileges.
- You supply a user ID, certificate, and passphrase. You create the certificate and use that certificate when you connect to the agent. This certificate must be accessible from the server.

Windows SMB protocol

You supply the user ID and password. The user ID must have administrative privileges. You must also enable **File & Printer Sharing for Microsoft Windows** under **Network Properties** for the connected network adapter. Make sure that the Server service is running on the system.

RExec protocol

You supply the user ID and password. The user ID must have administrative privileges. The system must be enabled for remote execution of commands.

RSH protocol

You supply the user ID with administrative privileges. The system must be enabled for the user to be able to run commands through the remote shell.

To enable RSH, edit the .rhosts file in the login directory of the user. The rhosts file provides access for the user to connect from a remote system.

If you elect to use a Microsoft directory, you can also get a list of domain computers. You are required to enter the domain controller, user ID, password, and certificate location. After the list is displayed, you can select a list of computers on which to deploy the agent.

After the agent is deployed, a probe job is automatically run.

The agent that is deployed on a Windows system has its registry information in the Windows registry. For UNIX or Linux®, the registry information is stored in the following file: `/etc/Tivoli/TSRM/registryNA`.

Planning for Internet Protocol Version 6

IBM Spectrum® Control supports Internet Protocol Version 6 (IPv6) for communication between its components. The key IPv6 enhancement is the expansion of IP address spaces from 32 bits (up to 15 characters in length) to 128 bits (up to 45 characters in length).

Overview

You can install and run IBM Spectrum Control on systems that are enabled for IPv4, IPv6, or dual stack. *Dual stack* indicates that a system has both the IPv4 and IPv6 stacks enabled and both addresses configured.

You can use IPv6 addresses if one of the following conditions are met:

- The system where IBM Spectrum Control is installed is IPv6.
- Dual stack (IPv4 and IPv6) is enabled.

IBM Spectrum Control can communicate with the following external components over IPv6:

- SMI-S CIM agents
- SLP directory agents
- SNMP agents
- Storage resources
- SMTP server

Note: These external components must be IPv6 enabled to communicate with IBM Spectrum Control by using the IPv6 protocol.

Not all components and their related functions of IBM Spectrum Control are enabled for IPv6. Any functions that are not enabled for IPv6 are unavailable through the user interface when you install IBM Spectrum Control on an IPv6-only system.

If you have a system that is configured for dual stack networking, IBM Spectrum Control defaults to IPv4 addressing.

The preferred IPv6 address representation is written as eight groups of four hexadecimal digits `xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`, where each x is a hexadecimal that represents 4 bits. You can also specify IPv6 addresses by using shortened formats that omit leading zeros or use double colons in place of a series of zeros. You can specify only one double colon (::) in an IPv6 address.

The following examples show an IPv6 address in the long form and in the short form:

Long form: `2001:db8:0:0:0:0:0:0`

Short form: `2001:db8::`

IPv6 and IBM Spectrum Control installation

If you are installing IBM Spectrum Control in an IPv6-only environment, the following requirements apply:

- For Windows operating systems, the IPv4 loopback interface is enabled.
 - For AIX® or Linux® operating systems, the localhost option resolves to the IPv4 address 127.0.0.1, instead of the IPv6 address ::1. This value is specified in the `/etc/hosts` file.
- For Windows operating systems, the localhost option is specified in the `C:\Windows\system32\drivers\etc\hosts` directory and file.

Table 1. IPv4 and IPv6 configurations in a multiple-servers environment

If the database repository server is configured for...	The IBM Spectrum Control server is configured for...
IPv4	IPv4
IPv4	Dual stack
Dual stack	IPv4
Dual stack	Dual stack
Dual stack	IPv6
IPv6	IPv6

Restriction: When you install the product on servers that are configured for IPv6 only, you cannot use IBM® Tivoli Enterprise Console® events as triggered actions in alerts.

Planning to use LDAP for IBM Spectrum Control authentication

The Lightweight Directory Access Protocol (LDAP) is an application protocol that you can use to query and modify directory services running over TCP/IP. The IBM Spectrum® Control installation program establishes a default authentication configuration using the federated repositories feature of the IBM® WebSphere® Application Server Liberty. You can configure IBM Spectrum Control for LDAP authentication as a post-installation activity.

About this task

A *directory* is a set of objects with similar attributes that are organized in a logical and hierarchical manner. An LDAP directory tree often reflects various political, geographic, and organizational boundaries, depending on the model chosen. The directory might contain entries representing people, organizational units, printers, documents, groups of people, or anything else that represents a tree entry (or multiple entries).

In the federated repositories framework, the IBM Spectrum Control installation program creates the following repositories:

File-based user repository

This repository contains the `tpcFileRegistryUser` user ID. This user password is the same as the common user password that you enter during the IBM Spectrum Control installation.

Operating system repository

This repository contains the users and groups managed by the local operating system.

With this default authentication configuration, you cannot use the IBM Spectrum Control single sign-on feature. Storage system element managers do not support the operating system repository or the file-based repository for single sign-on, even if the element manager is installed on the same system as IBM Spectrum Control.

To change the user authentication configuration, add or remove an LDAP repository in the federated repositories framework.

To authenticate user names on the IBM Spectrum Control by using an instance of Microsoft Active Directory, complete one of the following options:

Authenticate User names: Method one

1. Install IBM Spectrum Control on an AIX®, Linux®, or Windows operating system.
2. Configure IBM Spectrum Control for LDAP authentication.

The target LDAP repository is an instance of Active Directory running on Windows 2003 or later. This method provides the following benefits:

- The IBM Spectrum Control target system does not have to be a Windows system.
- The IBM Spectrum Control target system does not have to be a member of a Windows domain.
- There are fewer configuration steps for IBM Spectrum Control or Windows.
- The IBM Spectrum Control user names and group names do not have to include the Windows Domain Name.
- You can use the single-sign on feature between IBM Spectrum Control and storage system element managers.

Authenticate User names: Method two

1. Install IBM Spectrum Control on a Windows system that is a member of a Windows domain or on the Windows Domain Controller.
2. Configure IBM Spectrum Control for operating system authentication.

Planning for storage management

Plan for how to use IBM Spectrum® Control to manage the storage and storage resources in your environment.

- [Planning for storage systems](#)
To effectively plan your storage subsystems, you must understand the terminology.
- [Planning to monitor performance](#)
IBM Spectrum Control can collect performance metrics for resources that use the native interfaces, such as DS8000®, XIV®, SAN Volume Controller, Storwize® V7000 Unified, or Storwize V7000 storage systems, or for storage systems and Fibre Channel switches that use SMI-S providers (also called CIM agents or CIMOMs). The resources that use SMI-S providers must be SMI-S 1.1 compliant. To collect performance metrics for IBM FlashSystem® 900, the SNMP agent must be enabled on the storage system.
- [Planning for switches and fabrics](#)
IBM Spectrum Control supports IBM®, Brocade, and Cisco switches. Use IBM Spectrum Control to help manage the switches and fabrics that connect host systems and applications to storage resources.
- [Planning for VMware](#)
IBM Spectrum Control supports the VMware vSphere components ESXi and vCenter Server.
- [Planning for files systems and volume managers](#)
Information about the file system formats and volume managers that are supported by IBM Spectrum Control can help you set up and configure your environment.
- [Planning for PowerHASystemMirror for AIX](#)
IBM Spectrum Control supports Storage Resource agents that are installed on IBM PowerHA® SystemMirror® for AIX® nodes. Use this information to configure the

PowerHA SystemMirror for AIX environment before you use it with IBM Spectrum Control.

- [NAS support](#)

You can use the information to plan your Network Attached Storage (NAS) support.

- [Microsoft Cluster Server](#)

IBM Spectrum Control can monitor and report on Microsoft Cluster Server (MSCS) clustered nodes and cluster resource groups.

- [Planning for the Virtual I/O Server](#)

You can use the Storage Resource agent to gather information about Virtual I/O Servers. Before you can monitor Virtual I/O Servers, you must plan on how to install the agents in your environment.

- [Planning to monitor Db2](#)

It is a good practice to monitor your IBM Db2® environment to better understand what is happening inside your Db2 data server. Db2 11.5 or later includes enhancements that make monitoring Db2 database environments more comprehensive with higher granularity of control.

Planning for storage systems

To effectively plan your storage subsystems, you must understand the terminology.

About this task

IBM Spectrum® Control uses abstract terminology, such as pools and volumes, to help you view and manage your heterogeneous storage systems. Most of these terms are derived from SMI-S, which already provides a common model for storage subsystems.

storage pool

A collection of storage capacity that provides the capacity requirements for a volume. A pool has certain storage capabilities, which indicate the range of quality of service requirements that can be applied to objects created from the pool.

primordial pool

A type of storage pool. This pool might contain unformatted, unprepared, or unassigned capacity. Storage capacity is drawn from the primordial storage pool to create concrete storage pools. The primordial storage pool aggregates storage capacity that has not been assigned to a concrete storage pool. Storage volumes are allocated from concrete storage pools. For DS8000® storage systems, the primordial pool is the disk groups or array sites that are installed in the machine but have not yet been configured into RAID arrays. Primordial pools for SAN Volume Controller or Storwize® V7000 are those MDisk that are available, but have not yet been configured to any MDisk Group (one primordial pool per back-end controller). The Primordial pool for the XIV® is a virtual concept that represents the aggregation of system-wide deallocated storage capacity that is available but unassigned to the XIV storage pools.

storage volumes

Allocations of storage capacity that is exposed from a system through an external interface. In SCSI terms, these storage volumes are logical units.

The following table shows the mapping of the IBM Spectrum Control terms to the device-specific terms.

IBM Spectrum Control terms:	Storage pool	Primordial pool	Storage volume	Disk
Device terms				
DS8000	Extent pool	Unconfigured disk groups or array sites	Volume	Disk drive module (DDM)
SAN Volume Controller or Storwize V7000 or FlashSystem 7200	MDisk group	Managed disks	Volume	MDisk
XIV	Storage pool	not applicable	Volume	Disk
VMAX	Data Pool	not applicable	Volume	Disk
VNX and VNXe	Pool/RAID Group	not applicable	LUN	Disk

Many disk arrays provide an interface for the administrator to specify which initiators can access what volumes through which target ports. The effect is that the volume is only visible to SCSI commands that originate from the specified initiators through specific sets of target ports. There might also be a capability to select the SCSI Logical Unit Number as seen by an initiator through a specific set of ports. The ability to limit access is called *device masking*. The ability to specify the device address that is seen by particular initiators is called device mapping. For SCSI systems, these terms are known as LUN masking and LUN mapping. In IBM Spectrum Control, masking and mapping are handled through host assignment.

Note:

- DS8000 storage systems track the number of spares still available after a spare disk has been used to replace a failing disk drive in the device by marking the failing drive with an operational status of "Predictive Failure." If there are not enough remaining spare disks for a rank, then the operational status of the disk also shows a status of "Error."
If a DS8000 storage system marks a disk drive with a status of "Predictive Failure" but not "Error", then IBM Spectrum Control shows a green icon with a consolidated status of "OK" in the health overlay, and does not create an alert.

IBM Spectrum Control supports IBM® and independent storage vendor systems through the native interfaces or with Storage Management Interface Specification (SMI-S) compatible interfaces. This support includes storage provisioning, asset reporting, and capacity reporting.

The following table shows the IBM storage systems that you can use, and the data sources that you can use with them:

IBM storage system	Data source
SAN Volume Controller	Native interface
DS8000 series	Native interface
XIV	Native interface
IBM Spectrum Accelerate	Native interface. Storage provisioning is not supported.
IBM FlashSystem® devices that run IBM Spectrum Virtualize	Native interface
IBM FlashSystem 900	Native interface. Performance monitoring and storage provisioning are not supported by the native interface. You must use SNMP to collect performance data.
FlashSystem A9000	Native interface
FlashSystem A9000R	Native interface
Storwize V7000	Native interface

IBM storage system	Data source
Storwize V7000 Unified	Native interface
IBM Spectrum Scale	Native interface. Storage provisioning is not supported.
IBM Cloud Object Storage	Native interface. Storage provisioning is not supported.
SMI-S certified storage systems	For information about SMI-S certified storage systems, see https://www.snia.org/ctp/conforming_providers_archive .

Restriction:

- A specific SAN Volume Controller system must not be managed by more than one IBM Spectrum Control server at the same time.

For the most current information about storage systems, firmware, and provider levels that are supported, see the IBM Spectrum Control support website at <https://www.ibm.com/support/pages/node/388393>

- **[Planning for monitoring storage systems](#)**
Plan for monitoring the storage systems in your environment.
- **[Planning for SMI-S providers](#)**
SMI-S providers are provided by the vendor of certain types of storage systems and switches. IBM Spectrum Control communicates with an SMI-S provider to collect information about the resources that the agent manages.
- **[Planning for TagmaStore CIM agents](#)**
The TagmaStore CIM agents are provided by Hitachi Data Systems for the TagmaStore storage system. The TagmaStore CIM agent collects information from the TagmaStore storage system.
- **[Planning for the native interface](#)**
IBM Spectrum Control communicates with DS8000, the SAN Volume Controller, XIV, IBM Spectrum Accelerate, and Storwize V7000 storage systems through their native interfaces. You do not need to install and maintain Common Information Model (CIM) agents to collect data from these storage systems.
- **[Storage capacity of volumes](#)**
IBM Spectrum Control displays the storage capacity of volumes.
- **[Planning for DS8000](#)**
Use this information to plan for using the IBM DS8000 storage system with IBM Spectrum Control.
- **[Planning for SAN Volume Controller](#)**
Plan for using the SAN Volume Controller with IBM Spectrum Control.
- **[Planning for IBM Spectrum Virtualize for Public Cloud](#)**
IBM Spectrum Virtualize is a software-defined storage solution that has been proven for years in SAN Volume Controller and the IBM Storwize family. IBM Spectrum Virtualize for Public Cloud extends that solution to a hybrid-cloud or cloud-based model, where servers, storage, and network infrastructure are delivered in a public cloud environment. It can be deployed on either IBM® Cloud or Amazon Web Services (AWS) cloud infrastructures.
- **[Planning for Storwize V7000](#)**
Storwize V7000 is a hardware and software solution that provides unmatched performance, availability, advanced functions, and highly scalable capacity.
- **[Planning for Storwize V7000 Unified](#)**
The IBM Storwize V7000 Unified system is a virtualizing redundant array of independent disks (RAID) storage system that supports both block protocols and file protocols. This unified system includes Storwize V7000 File Module and the Storwize V7000 storage system.
- **[Planning for XIV storage systems](#)**
Plan to monitor XIV storage systems with IBM Spectrum Control.
- **[Planning for Dell EMC storage systems](#)**
Dell EMC resources provide block and unified storage for organizations with network-attached storage or storage area network environments that have file and block-level services. IBM Spectrum Control provides enhanced monitoring for Unity, VMAX, VNX, and VNXe storage systems.
- **[Planning for Hitachi storage systems](#)**
Hitachi resources provide block storage for organizations with network-attached storage or storage area network environments that have block-level services. IBM Spectrum Control provides enhanced monitoring for Hitachi VSP F and G Series storage systems.
- **[Planning for NetApp device support](#)**
NetApp devices provide unified storage for organizations with network-attached storage or storage area network environments that have file and block-level services. Plan for how to use IBM Spectrum Control to monitor your NetApp storage systems.
- **[Planning for Pure Storage systems](#)**
Pure storage resources provide block storage for organizations with network-attached storage or storage area network environments that have block-level services. IBM Spectrum Control provides enhanced monitoring for Pure FlashArray//M and FlashArray//X storage systems.

Planning for monitoring storage systems

Plan for monitoring the storage systems in your environment.

Procedure

1. Prepare for collecting data about storage systems.
For storage systems that require SMI-S providers (also called CIM agents or CIMOMs), ensure that the following conditions are met before you add the SMI-S provider to IBM Spectrum® Control:
 - The version of SMI-S provider and firmware for the storage system is supported. For information about the SMI-S provider and firmware that is supported, see [IBM Spectrum Control interoperability matrix for storage systems](#).
 - An SMI-S provider is installed on a different server than the IBM Spectrum Control server.
 - For storage systems on a private network, ensure that the SMI-S provider is installed on a gateway machine so that the IBM Spectrum Control server can communicate with that agent.
 - The SMI-S provider is configured to manage the intended storage system.
 For information about how to install and configure the SMI-S provider, contact the vendor of the SMI-S provider.
2. Add storage systems for monitoring with the IBM Spectrum Control GUI.
When you add storage systems for monitoring, you can automatically schedule the following data collection jobs:
 - Probes are data collection jobs that collect asset, status, and storage data about storage systems.
 - Performance monitors are data collection jobs that collect performance information about storage systems.
 Tips for probes and SMI-S providers: Running a probe puts additional workload on the IBM Spectrum Control server, the repository database, and the SMI-S providers. For probes, the additional workload is especially important if multiple devices are managed by the same SMI-S provider and are probed concurrently. As

a result, the probe completion time is impacted by the number of devices that are probed in parallel. This fact is especially important if the SMI-S provider machines are lower-end machines of the hardware prerequisites.

To help keep the workload low, run probes at a time when IBM Spectrum Control and the network are not used heavily, which is typically at night. To allow this workload distribution, schedule probes for each storage system to run one after the other.

The following example shows how you might schedule probes. This environment assumes that you have four storage systems (SS1, SS2, SS3, and SS4), with three of them having a probe duration of less than 30 minutes and one (SS3) taking 90 minutes to complete. You can schedule the probes according to the following table:

Probe	Probe 1	Probe 2	Probe 3	Probe 4
Storage system	SS1	SS2	SS3	SS4
Start time	0:30	1:30	2:30	4:30

Probes collect detailed information about the configuration of the storage systems and the properties that describe those configuration elements. Information about the following elements is gathered:

- Component computer systems (for SAN Volume Controller or Storwize® V7000 nodes)
- Pools
- Volumes
- Fibre Channel ports
- Disks
- Host-to-volume assignments
- Relationships among the internal resources (for example, pool-to-volume relationships)

Probe storage systems once every one or two days. The following factors can influence this schedule:

- Is the storage system managed by IBM Spectrum Control or is a different tool used (for example, a device-specific CLI)?
- How extensively are IBM Spectrum Control alerts used?

3. Create alert definitions to determine when and how you are alerted to conditions or violations on storage systems. Alerts are triggered by conditions that are detected during data collection and event processing. Use alert policies to manage the alert definitions that apply to groups of resources of the same type. You can use default alert policies, or create your own by copying a default policy.

For some storage systems such as IBM Spectrum Accelerate and the XIV®, events are polled every minute from the resource. For IBM Spectrum Scale, status change events are polled frequently, typically within minutes. For other resources, events are subscription-based, where the resource itself or a data source such as an SMI-S provider sends the events to IBM Spectrum Control when conditions change on the resource.

Examples of storage systems that use subscription-based event processing include SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, and FlashSystem V9000. For these storage systems, a probe is automatically run when many events are received from the storage system in a short time period. To avoid performance bottlenecks, probes are run only every 20 minutes.

4. Configure data retention.

On the History Retention page in the IBM Spectrum Control GUI, configure how long to retain data that is collected about storage systems. By configuring data retention, you can control the amount of data that is retained and available for historical analysis and charting. The longer you keep the data, the more informative your analysis.

Planning for SMI-S providers

SMI-S providers are provided by the vendor of certain types of storage systems and switches. IBM Spectrum® Control communicates with an SMI-S provider to collect information about the resources that the agent manages.

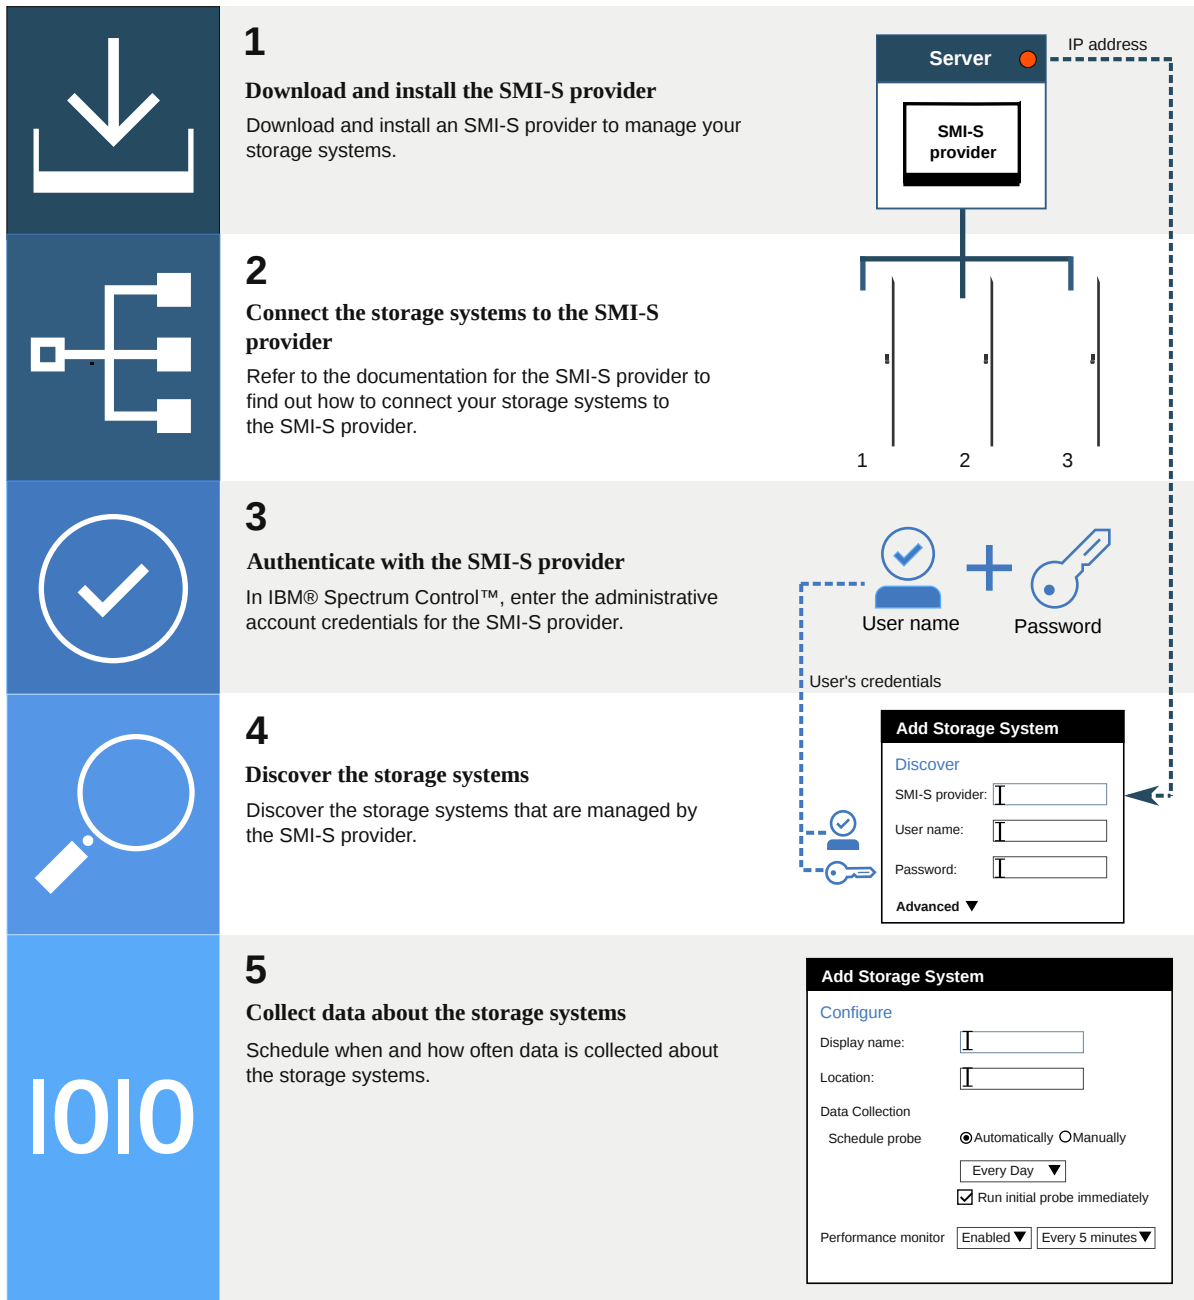
About this task

For some storage systems, SMI-S providers are required for collecting storage asset information, provisioning, alerting, and performance monitoring. SMI-S providers must conform to the SNIA SMI-S specification to provide a communication transport between IBM Spectrum Control and storage resources.

Note: For DS8000®, the XIV®, IBM Spectrum Accelerate, IBM FlashSystem® devices that run IBM Spectrum Virtualize, SAN Volume Controller, Storwize® V7000 Unified, Storwize V7000, IBM Spectrum Scale, and IBM® Cloud Object Storage, you do not need an SMI-S provider. IBM Spectrum Control communicates with these storage systems by using the native interface.

The SMI-S providers can be referred to by various names, such as CIM agent, CIMOM (CIM Object Manager) agent, or SMI-S agent. A CIM agent consists of a CIMOM and an SMI-S device provider for the managed resource. The SMI-S provider can be a separate agent installation or can be embedded in the resource itself, as is the case with Cisco fabric switches. In this case, there is no proxy agent to install and IBM Spectrum Control is configured to point to the managed resource itself.

After the SMI-S provider is installed and configured, IBM Spectrum Control can be configured to communicate with it.



SMI-S providers are not required for storage systems that communicate through a native interface. For those resources that require SMI-S providers to be set up, review the SMI-S provider documentation to understand how many resources that the SMI-S provider can be configured to manage. The memory that is consumed by the SMI-S provider includes the memory that is required for IBM Spectrum Control to probe a resource through the SMI-S provider. The SMI-S provider is typically affected by the number of resources that the SMI-S provider manages. If no guidelines are available in the SMI-S provider documentation for the storage system, limit the number of storage systems to three per SMI-S provider.

For information about the latest SMI-S provider support for IBM Spectrum Control, see the <https://www.ibm.com/support/pages/node/388393> site and go to the *Switches and Directors* and *Storage* sections.

For information about certified SMI-S devices, see http://www.snia.org/ctp/conforming_providers.

Related reference

- [Performance metrics for other storage systems](#)

Planning for TagmaStore CIM agents

The TagmaStore CIM agents are provided by Hitachi Data Systems for the TagmaStore storage system. The TagmaStore CIM agent collects information from the TagmaStore storage system.

IBM Spectrum® Control supports Hitachi Device Manager software.

All volumes are created from a storage pool that is allocated from a primordial storage pool and an imported primordial pool. A volume cannot be created over both local and virtual extents.

For LUN and volume correlation, the host machine must have the Storage Resource agent installed and the TagmaStore resource must be in the same SAN fabric. You must also have a zone configured in the active zone set between the ports of the host machine and the ports of the TagmaStore device. The Storage Resource agent must monitor the fabric to which the host is connected.

For back-end correlation, the TagmaStore device ports and back-end storage system ports must be in the same zone and the back-end storage system has assigned storage volumes to all ports of the TagmaStore device.

The following TagmaStore storage system events are reported to IBM Spectrum Control by CIM agents:

- Generation of a volume
- Deletion of a volume
- Allocation of a path
- Cancellation of a path

When you delete a volume, you can only delete the latest created volume. For example, if you create five volumes, you need to delete the volumes in reverse order from the creation order.

General procedure

Complete the following general steps to monitor a TagmaStore storage system in IBM Spectrum Control:

1. In IBM Spectrum Control, add TagmaStore storage system for monitoring. To add the storage system, you must specify information about the CIM agent that manages it.
2. Run a probe to collect information about the storage system.
3. Go to the Block Storage Systems page, right-click the TagmaStore storage system, and select View Details.

Related information

- <https://www.ibm.com/support/pages/node/388393>
- <http://www.ibm.com/support/docview.wss?uid=swg21667294>

Planning for the native interface

IBM Spectrum® Control communicates with DS8000®, the SAN Volume Controller, XIV®, IBM Spectrum Accelerate, and Storwize® V7000 storage systems through their native interfaces. You do not need to install and maintain Common Information Model (CIM) agents to collect data from these storage systems.

Support for the native interfaces

For the DS8000, IBM Spectrum Control communication with the storage system is through the ESSNI interface. Through this interface, IBM Spectrum Control collects device information, performance information, and can detect changes and complete provisioning operations. This support simplifies the IBM Spectrum Control configuration and the launch in context configuration required to manage the storage system. IBM Spectrum Control also supports dual HMC environments.

For the SAN Volume Controller or Storwize V7000, IBM Spectrum Control uses the CLI interface to communicate with the storage system. Through this interface, IBM Spectrum Control collects device information, performance information, and can detect changes and complete provisioning operations. This support simplifies the IBM Spectrum Control configuration and the launch in context configuration required to manage the storage system.

For the XIV or IBM Spectrum Accelerate, IBM Spectrum Control uses the XIV XML API interface to communicate with the storage system. The IBM Spectrum Control server communicates directly with the storage system using XML over an SSL socket. Through this interface, IBM Spectrum Control collects device information, performance information, and can detect changes and complete provisioning operations. This support simplifies the IBM Spectrum Control configuration and the launch in context configuration required to manage the storage system.

Use the add storage system action on storage system pages to define native interface connection information as described in [Adding resources](#).

Storage capacity of volumes

IBM Spectrum® Control displays the storage capacity of volumes.

About this task

The disk storage that IBM Spectrum Control supports is expressed in powers of two:

1 KB = 2 to the power of 10 bytes (1024)
1 MB = 2 to the power of 20 bytes (1 048 576)
1 GB = 2 to the power of 30 bytes (1 073 741 824)
1 TB = 2 to the power of 40 bytes (1.09951E+12)
1 PB = 2 to the power of 50 bytes (1.1259E+15)

This matches the convention of many storage systems but does not match the behavior of the the IBM® XIV® Storage System or IBM Spectrum Accelerate. Therefore, a volume that is created by using the the XIV GUI with a size of 17 GB might be displayed by IBM Spectrum Control as 16 GB in size. The DS8000® storage system GUIs display the storage as powers of two.

Planning for DS8000

Use this information to plan for using the IBM® DS8000® storage system with IBM Spectrum® Control.

IBM System Storage® DS8000 series is a high-performance, high-capacity series of disk storage that is designed to support continuous operations. The DS8000 uses the IBM Spectrum Control native interface for connecting to the device. CIM agents are no longer required for the DS8000.

For more information about DS8000, see the IBM Docs at <https://www.ibm.com/docs/en/ds8870>.

To view information about a DS8000, complete the following steps:

1. In the menu bar, go to Storage > Block Storage Systems.
2. Click Add Storage System.
3. Click the DS8000 icon.
4. Add the storage system for monitoring. A probe automatically collects configuration information about the storage system.
5. View information about the DS8000 storage system in the IBM Spectrum Control GUI.

Important information about host connections: For DS8000 storage systems, if host connections were created by using the **mkhostconnect** command, those host connections might not be displayed on the details page for the storage system in IBM Spectrum Control. Hosts that were configured with the **mkhost** command are displayed properly.

To ensure that the host connections that were configured with **mkhostconnect** and any related agentless servers are displayed properly in IBM Spectrum Control, complete the steps at [Host connections for DS8000 storage systems are not being displayed](#).

Related reference

- [Performance metrics for DS8000](#)

Planning for SAN Volume Controller

Plan for using the SAN Volume Controller with IBM Spectrum® Control.

When you have hosts connected to storage subsystems which have multi-pathing enabled, you must have the multi-pathing subsystem device drivers (SDD) installed on the hosts. For more information, see the documentation for SAN Volume Controller at <https://www.ibm.com/docs/en/sanvolumecontroller/8.3.1> or <https://www.ibm.com/docs/en/sanvolumecontroller/8.4.0>.

SAN Volume Controller 6.1 or later

SAN Volume Controller includes IBM® System Storage® Easy Tier®, a function that responds to the presence of solid-state drives (SSDs) in a storage pool that also contains hard disk drives (HDDs). The system automatically and nondisruptively moves frequently accessed data from HDD MDisks to SSD MDisks, thus placing such data in a faster tier of storage.

Easy Tier eliminates manual intervention when assigning highly active data on volumes to faster responding storage. In this dynamically tiered environment, data movement is seamless to the host application regardless of the storage tier in which the data resides.

SAN Volume Controller supports these tiers:

Generic SSD tier

The SSD tier exists when SSDs are in the storage pool. The SSDs provide greater performance than hard disk drives (HDDs).

Generic HDD tier

The HDD tier exists when HDDs are in the storage pool.

All MDisks belong to one tier or the other, which includes MDisks that are not yet part of a storage pool.

If you create a storage pool (managed disk group) with both generic SSD MDisks and generic HDD MDisks, Easy Tier is automatically turned on for pools with both SSD MDisks and HDD MDisks. SAN Volume Controller does not automatically identify external SSD MDisks; all external MDisks are put into the HDD tier by default. You must manually identify external SSD MDisks and change their tiers.

Note:

- The terminology for this release of SAN Volume Controller has changed:
 - Error is now event
 - MDisk is now storage pool
 - Space-efficient is now thin provisioning
 - VDisk is now volumeHowever, the terms have not been changed in IBM Spectrum Control reports and views.
- In SAN Volume Controller 5.1 or earlier, there was a one-to-one mapping between MDisks and internal SSD drives. For SAN Volume Controller 6.1 or later, there are now one-to-many mapping between MDisks and drives.

In addition to information already available for SAN Volume Controller, IBM Spectrum Control provides information about the tier, tier capacity, and tier free capacity.

You can also open the management GUI for SAN Volume Controller from inside the IBM Spectrum Control GUI.

Upgrading SAN Volume Controller 5.1

Currently SAN Volume Controller Version 6.1 does not support the solid-state drives. This support will be added in a future fix pack.

If you are upgrading SAN Volume Controller 5.1 to 6.1, you have the following options:

- Do not upgrade SAN Volume Controller 5.1 to 6.1 until the solid-state drive support is available.
- Follow these steps to upgrade SAN Volume Controller 5.1:
 1. Migrate data off the solid-state drive MDisks.
 2. Unconfigure (unmanage) the solid-state drive MDisks.
 3. Run a probe to collect the latest configuration data about the SAN Volume Controller.
 4. Upgrade the SAN Volume Controller 5.1 to 6.1.

Note: SAN Volume Controller does not permit an upgrade of SAN Volume Controller 5.1 to 6.1 while there are managed solid-state drive MDisks. After the upgrade, the unmanaged solid-state drive MDisk objects are marked as unused drive objects. You cannot configure these drive objects into a RAID.

Related reference

- [Performance metrics for resources that run IBM Spectrum Virtualize](#)

Planning for IBM Spectrum Virtualize for Public Cloud

IBM Spectrum Virtualize is a software-defined storage solution that has been proven for years in SAN Volume Controller and the IBM® Storwize® family. IBM Spectrum Virtualize for Public Cloud extends that solution to a hybrid-cloud or cloud-based model, where servers, storage, and network infrastructure are delivered in a public cloud environment. It can be deployed on either IBM® Cloud or Amazon Web Services (AWS) cloud infrastructures.

With IBM Spectrum® Control, you can view the capacity, space usage, and performance of your IBM Spectrum Virtualize for Public Cloud storage systems. Other monitoring features, such as alerting, health checking, advanced analytics, and reporting are also supported.

Benefits

IBM Spectrum Control can help you predict and prevent storage problems before they impact your business. Here are some key benefits of using IBM Spectrum Control to monitor your IBM Spectrum Virtualize for Public Cloud storage systems:

- View detailed information about capacity, storage usage, and performance.
- Monitor health, status, and availability.
- Use alerts and alert policies to be notified of conditions and potential problems.
- Use advanced analytics to provision, optimize, and reclaim storage.
- Create and share reports about inventory, capacity, performance, and storage consumption.

How to monitor

Before you can add an IBM Spectrum Virtualize for Public Cloud storage system for monitoring, you must ensure that IBM Spectrum Control can connect to it. To enable a connection, you can use the site-to-site VPN IPsec tunnel that exists between the on-premises environment and the IBM Spectrum Virtualize for Public Cloud instances in AWS or IBM Cloud®.

Tips:

- For more information about configuring the connection between IBM Spectrum Control and IBM Spectrum Virtualize for Public Cloud, see [Configuring IBM Spectrum Virtualize for Public Cloud for monitoring](#).
- IBM Spectrum Control uses the following ports for communicating with IBM Spectrum Virtualize for Public Cloud:
 - Outbound: 22 for SSH
 - Outbound: 5989. Optionally, the SSH key can be uploaded once at setup time using the 5989 port.

By default, IBM Spectrum Control creates a connection between IBM Spectrum Virtualize for Public Cloud and a local port to listen for events. The default range is 49152 to 65535. These ports are on the IBM Spectrum Control Device server only; you do not need to configure the firewall for them. The ports are used for a SSH tunnel to connect to the event port on IBM Spectrum Control. All communication to IBM Spectrum Control runs through the SSH connection. No other ports need to be open for events.

Supported features

View a detailed list of the features in IBM Spectrum Control that you can use to monitor IBM Spectrum Virtualize for Public Cloud storage systems:

Table 1. Supported features for monitoring IBM Spectrum Virtualize for Public Cloud

Resource Monitoring	Features	Supported
Understanding the environment	Monitor storage inventory and configuration. Includes information about type, model, serial number, and firmware.	✓
	Understand storage relationships, from volume and share down to server and application.	✓
	Explore virtualization relationships.	✓
	Explore replication relationships.	✓
	View dashboards to get insights into key aspects of your storage at a glance and one-click access to web-based element managers.	✓
Monitoring capacity	Collect storage consumption and capacity metrics.	✓
	View data reduction information.	✓
	View copy data information.	✓
	View internal storage tiers such as EasyTier.	✓
	Monitor the storage consumed by applications.	✓
Monitoring performance	Collect performance metrics about the workload on resources.	✓*
	View calculated metrics to gain insights into performance conditions.	✓
	Export performance data to a compressed file.	✓
	Drill down performance workflows to troubleshoot bottlenecks.	✓

Resource Monitoring	Features	Supported
	Compare the performance of resources.	✓
Monitoring health	Understand the health of resources.	✓
	Receive notifications when the status of a resource changes.	✓
	View the status of elements that are not represented as resources.	✓
Alerting	Alert on conditions within your storage environment.	✓
	Define alerts to identify issues based on multiple conditions.	✓
	Define alert policies to be notified of changes across related resources.	✓
Reporting	View predefined reports.	✓
	View chargeback and consumer reports.	✓
	Create custom reports by using the REST API.	✓
	Create rollup reports to view information across multiple IBM Spectrum Control servers.	✓
Analytics	Analyze business impact (applications, departments, and groups).	✓
	Optimize data placement with tiering.	✓
	Optimize capacity with reclamation.	✓
* Performance metadata for managed disks in IBM Spectrum Virtualize for Public Cloud is not available.		

For your reference: Want to learn more about IBM Spectrum Virtualize for Public Cloud? No problem. IBM provides the following documentation for your reference:

- [IBM Spectrum Virtualize for Public Cloud Knowledge documentation](#)
- [Redbook \[PDF\]: Implementation guide for IBM Spectrum Virtualize for Public Cloud 8.3](#)
- [Redbook \[PDF\]: IBM Spectrum Virtualize for Public Cloud on AWS Implementation Guide](#)

What's next

After you establish the method for connecting IBM Spectrum Control to a IBM Spectrum Virtualize for Public Cloud storage system, you can add that storage system for monitoring, alerting, and reporting.

For instructions on how to add storage systems for monitoring, see [Adding storage systems](#).

Related reference

- [Ports used by IBM Spectrum Control](#)
- [Required user roles for monitoring resources](#)
- [Performance metrics for resources that run IBM Spectrum Virtualize](#)

Planning for Storwize V7000

Storwize® V7000 is a hardware and software solution that provides unmatched performance, availability, advanced functions, and highly scalable capacity.

General concepts

Storwize V7000 offers IBM® storage virtualization, SSD optimization and “thin provisioning” technologies built in to improve storage utilization. The storage system can be reconfigured to meet changing needs quickly and easily. This solution helps to reduce costs without performance degradation.

The Storwize V7000 hardware consists of a set of drive enclosures. Control enclosures contain disk drives and two node canisters. The two nodes within the canisters make an I/O group that is attached to the SAN fabric. A single pair of nodes is responsible for serving I/O on a given volume. Because a volume is served by two nodes, there is no loss of availability if one node fails or is taken offline.

Storwize V7000 can be used as a traditional RAID storage system where the internal drives are configured into arrays, and volumes are created from those arrays. Storwize V7000 can also be used to virtualize other storage systems.

Storwize V7000 supports both regular and solid-state drives (SSDs). A Storwize V7000 system without any internal drives can be used as a storage virtualization solution.

Each Storwize V7000 node has two Ethernet ports that can be used for management. Ethernet port 1 must be configured with a management IP address and must be connected on all nodes in the system. The use of Ethernet port 2 is optional. At any point in time, only one node in the system can operate as the focal point for configuration and monitoring requests. This node is called the configuration node and is the only node that activates the management IP addresses.

Each Storwize V7000 can have zero to four management IP addresses. You can assign up to two IPv4 addresses and up to two IPv6 addresses.

All configuration, monitoring, and service tasks are performed at the cluster level (a cluster is a Storwize V7000 system that consists of two nodes).

Storwize V7000 includes IBM System Storage® Easy Tier®, a function that supports solid-state drives (SSDs) in a storage pool that also contains hard disk drives (HDDs). The system automatically and nondisruptively moves frequently accessed data from HDD MDisks to SSD MDisks, thus placing such data in a faster tier of storage.

Easy Tier eliminates manual intervention when assigning highly active data on volumes to faster responding storage. In this dynamically tiered environment, data movement is seamless to the host application regardless of the storage tier in which the data resides. Manual controls exist so that you can change the default behavior, for example, such as turning off Easy Tier on storage pools that have both types of MDisks.

Using IBM Spectrum Control to monitor Storwize V7000

IBM Spectrum® Control communicates with Storwize V7000 storage systems through a native interface. No CIM agents are required.

Use IBM Spectrum Control to monitor Storwize V7000 systems, including asset, status, and configuration information, performance monitoring, trending, storage optimization, and provisioning. IBM Spectrum Control also provides information about the tier, tier capacity, and tier free capacity.

Restriction: You can't use IBM Spectrum Control to monitor the SAS ports on a Storwize device. To monitor those ports, use the management GUI for the device.

Related reference

- [Performance metrics for resources that run IBM Spectrum Virtualize](#)

Planning for Storwize V7000 Unified

The IBM® Storwize® V7000 Unified system is a virtualizing redundant array of independent disks (RAID) storage system that supports both block protocols and file protocols. This unified system includes Storwize V7000 File Module and the Storwize V7000 storage system.

Using IBM Spectrum Control to monitor Storwize V7000 Unified

IBM Spectrum® Control communicates with Storwize V7000 Unified storage systems through a native interface. No CIM agents are required.

You can add a Storwize V7000 Unified storage system as a block storage system, a file storage system, or both. IBM Spectrum Control collects different information depending on how you add a Storwize V7000 Unified storage system, which includes information about capacity and configuration (block and file), performance (block only), and trending (block and file). IBM Spectrum Control reports also provide information about the tier, tier capacity, and tier free capacity for Storwize V7000 Unified.

Viewing Storwize V7000 Unified storage systems in the user interface

To view information about a Storwize V7000 Unified system, go to the Storage > Block Storage Systems or Storage > File Storage Systems page in the GUI. You can right-click a storage system and select View Details or View Properties to see additional information.

Related reference

- [Performance metrics for resources that run IBM Spectrum Virtualize](#)

Planning for XIV storage systems

Plan to monitor XIV® storage systems with IBM Spectrum® Control.

IBM® XIV Storage System is a disk storage architecture designed to eliminate the complexity of administration and management of storage. The XIV parallelized architecture of the system, optimal exploitation of all system components (including disks, CPUs, and switches), and unique caching architecture all translate into excellent performance.

The unique balancing of all data across system components prevents the occurrence of hot spots. With all components working under the same load, performance and reliability are exceptional.

XIV storage systems use large capacity Serial Advanced Technology Attachment (SATA) disk drives which optimizes the use of disk capacity, resulting in outstanding power consumption without compromising performance.

XIV storage systems are designed to be scalable in storage, interfaces, cache, CPU power, and internal bandwidth. The architecture supports each aspect to grow independently, resulting in a scalable system in both capacity and performance.

XIV storage systems provide data protection and availability. All disk drives, modules, switches, and uninterruptible power supply units are fully redundant, ensuring high reliability and excellent performance.

The built-in thin provisioning of XIV storage systems helps reduce direct and indirect costs by allowing users to install capacity only for data written. You can grow your data capacity over time with minimal management effort.

For more information about XIV storage systems, see the XIV documentation at <https://www.ibm.com/docs/en/xiv-storage-system>.

IBM Spectrum Control support

IBM Spectrum Control uses the native interface for XIV storage systems to collect metadata. Because of the dynamic XIV architecture, where volume data is spread over all disks, no data that describes the space allocation from volumes to a specific physical disk is available.

The XIV systems have the capability of defining up to three administrative nodes, each with their own IP address. In IBM Spectrum Control 5.1 or later, if an XIV is configured with multiple administrative nodes, IBM Spectrum Control detects the IP addresses for these nodes. If IBM Spectrum Control fails to connect to one of the IP addresses, then an attempt is made to connect to the XIV using one of the other IP addresses.

Restrictions:

- You cannot use IPv6 addresses when monitoring XIV storage systems.
- The password that IBM Spectrum Control uses to connect to an XIV storage system can't contain this special character: **&**.

Data collection for XIV systems

The following considerations apply to values shown in the XIV and IBM Spectrum Control GUI.

- IBM Spectrum Control displays capacity values in base 2 (GiB), while the XIV management GUI displays capacity values in base 10 (GB, TB). Even though different units of measurement are used, the storage values are equivalent. For more information about units of measurement, see [Units of measurement for storage data](#).

- The XIV differentiates thin provisioned pools from non-thin provisioned (regular) pools and volumes in a thin provisioned pool and a non-thin provisioned pool. When hard size is equal to soft size, the pool is represented as a regular pool, otherwise the pool is represented as thin provisioned. However, the XIV CLI and CIM do not distinguish between regular and thin provisioned pools and all pool volume names in IBM Spectrum Control are prefixed with an asterisk (*) to indicate that the volume is thin provisioned.

Related reference

- [Performance metrics for XIV, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R](#)

Planning for Dell EMC storage systems

Dell EMC resources provide block and unified storage for organizations with network-attached storage or storage area network environments that have file and block-level services. IBM Spectrum® Control provides enhanced monitoring for Unity, VMAX, VNX, and VNXe storage systems.

IBM Spectrum Control 5.3.5 or later supports the Unity storage system by connecting directly to the device.

IBM Spectrum Control version 5.2.13 or later supports versions of Dell EMC SMI-S Provider or Dell EMC Solutions Enabler that are compliant with SMI-S 1.6. Through SMI-S 1.6, the following storage systems are monitored:

- VMAX family
- VNX family
- VNXe family

For other Dell EMC storage systems, monitoring is through SMI-S 1.2.

Supported versions: To view the versions of Dell EMC storage systems that are supported in IBM Spectrum Control, go to the [Dell EMC support page](#).

Dell EMC Unity Device configuration

A Unity device can be added to IBM Spectrum Control as a block storage system, a file storage system, or both. When you add a Unity storage system, IBM Spectrum Control collects data by connecting directly to the storage system.

Antivirus software restriction: If your antivirus software is set on the maximum mode, it might prevent your Unity storage systems from being added to IBM Spectrum Control. For more information about how to configure your antivirus software, see [Installation checklists for IBM Spectrum Control](#).

Benefits

IBM Spectrum Control can help you predict and prevent storage problems before they impact your business. Here are some key benefits of using IBM Spectrum Control to monitor your Dell EMC storage systems:

- View detailed information about capacity, storage usage, and performance.
- Monitor health, status, and availability.
- Use alerts and alert policies to be notified of conditions and potential problems.
- Use advanced analytics to reclaim storage.
- Create and share reports about inventory, capacity, performance, and storage consumption.

Supported features

View a detailed list of the features in IBM Spectrum Control that you can use to monitor Dell EMC Unity storage systems:

Table 1. Supported features for Dell EMC Unity

Resource Monitoring	Features	Supported
Understanding the environment	Monitor storage inventory and configuration. Includes information about type, model, serial number, and firmware.	✓
	Understand storage relationships, from volume and share down to server and application.	✓
	Explore virtualization relationships.	
	Explore replication relationships.	✓
	View dashboards to get insights into key aspects of your storage at a glance and one-click access to web-based element managers.	✓
	Support for multiple protocols and storage types, such as FC, iSCSI, NVMe, CIFS, and NFS.	✓
Monitoring capacity	Collect storage consumption and capacity metrics.	✓
	View data reduction information.	✓
	View copy data information.	✓
	View internal storage tiers such as Easy Tier® and EMC FAST.	✓
	Monitor the storage consumed by applications.	✓
Monitoring performance	Collect performance metrics about the workload on resources.	✓
	View calculated metrics to gain insights into performance conditions.	✓
	Export performance data to a compressed file.	✓
	Drill down performance workflows to troubleshoot bottlenecks.	✓
	Compare the performance of resources.	✓
Monitoring health	Understand the health of resources.	✓
	Receive notifications when the status of a resource changes.	
	View the status of elements that are not represented as resources.	✓
Alerting	Alert on conditions within your storage environment.	✓

Resource Monitoring	Features	Supported
	Define alerts to identify issues based on multiple conditions.	✓
	Define alert policies to be notified of changes across related resources.	✓
Reporting	View predefined reports.	✓
	View chargeback and consumer reports.	✓
	Create custom reports by using the REST API.	✓
	Create rollup reports to view information across multiple IBM Spectrum Control servers.	✓
Analytics	Performance planning	✓
	Capacity planning	✓
	Business impact analysis (applications, departments, and groups)	✓
	Optimize data placement with tiering	
	Optimize capacity with reclamation.	✓

Restriction: In replication relationships, any snapshot volumes in either source or target storage systems that are not assigned to a host, are not stored in IBM Spectrum Control. The volumes are listed in the Copy Data page but do not have hyperlinks.

Metadata collection schedule

By default, asset, capacity, and configuration metadata is aggregated and collected daily. You can schedule daily capacity and inventory reports to gain insights about your Dell EMC storage systems.

For Dell EMC Unity storage systems, the collection intervals for performance metadata are 5 minutes and 60 minutes. The default interval is 5 minutes.

Performance metadata is collected every 15 minutes for VMAX, VNX, and VNXe block storage systems.

What's next

Before you can monitor, alert, and report on your Dell EMC storage systems, you must add them to IBM Spectrum Control.

For instructions on how to add storage systems for monitoring, see [Adding storage systems](#).

Related reference

- [Performance metrics for Dell EMC storage systems](#)

Planning for Hitachi storage systems

Hitachi resources provide block storage for organizations with network-attached storage or storage area network environments that have block-level services. IBM Spectrum® Control provides enhanced monitoring for Hitachi VSP F and G Series storage systems.

Hitachi VSP F and G Series Device configuration

A Hitachi VSP device can be added to IBM Spectrum Control as a block storage system. When you add a Hitachi VSP storage system, IBM Spectrum Control collects data by connecting to the Hitachi Command Suite that is managing the device. Connect to the Hitachi Device Manager to enable the Hitachi Export Tool to collect performance information from the storage system.

To view the versions of Hitachi VSP storage systems that are supported in IBM Spectrum Control, go to the [Hitachi support page](#).

Restrictions:

- Performance monitoring on Hitachi VSP storage systems is supported when IBM Spectrum Control is installed on Windows and Linux®. It is not supported when IBM Spectrum Control is installed on AIX®.
- You must set the Hitachi Performance Monitor Monitoring Switch to **Enabled** in the Hitachi Device Manager and then set the Sample Interval to **1 minute** or **5 minutes** only. Any other interval options are not supported. If performance monitoring is already enabled at an unsupported interval and you want to keep the previous monitoring data, first export the data by using the Export Tool and then restart monitoring.

Antivirus software restriction: If your antivirus software is set on the maximum mode, it might prevent your Hitachi VSP storage systems from being added to IBM Spectrum Control. For more information about how to configure your antivirus software, see [Installation checklists for IBM Spectrum Control](#).

Benefits

IBM Spectrum Control can help you predict and prevent storage problems before they impact your business. Here are some key benefits of using IBM Spectrum Control to monitor your Hitachi VSP storage systems:

- View detailed information about capacity, storage usage, and performance.
- Monitor health, status, and availability.
- Use alerts and alert policies to be notified of conditions and potential problems.
- Use advanced analytics to reclaim storage.
- Create and share reports about inventory, capacity, performance, and storage consumption.

Supported features

View a detailed list of the features in IBM Spectrum Control that you can use to monitor Hitachi VSP F and G Series storage systems:

Table 1. Supported features for Hitachi VSP F and G Series

Resource Monitoring	Features	Supported
Understanding the environment	Monitor storage inventory and configuration. Includes information about type, model, serial number, and firmware.	✓
	Understand storage relationships, from volume and share down to server and application.	✓
	Explore virtualization relationships.	
	Explore replication relationships.	✓
	View dashboards to get insights into key aspects of your storage at a glance and one-click access to web-based element managers.	✓
	Support for multiple protocols and storage types, such as FC, iSCSI, and PVMe.	✓
Monitoring capacity	Collect storage consumption and capacity metrics.	✓
	View data reduction information.	✓
	View copy data information.	✓
	View internal storage tiers such as Easy Tier®.	
	Monitor the storage consumed by applications.	✓
Monitoring performance	Collect performance metrics about the workload on resources.	✓
	View calculated metrics to gain insights into performance conditions.	✓
	Export performance data to a compressed file.	✓
	Drill down performance workflows to troubleshoot bottlenecks.	✓
	Compare the performance of resources.	✓
Monitoring health	Understand the health of resources.	✓
	Receive notifications when the status of a resource changes.	
	View the status of elements that are not represented as resources.	✓
Alerting	Alert on conditions within your storage environment.	✓
	Define alerts to identify issues based on multiple conditions.	✓
	Define alert policies to be notified of changes across related resources.	✓
Reporting	View predefined reports.	✓
	View chargeback and consumer reports.	✓
	Create custom reports by using the REST API.	✓
	Create rollup reports to view information across multiple IBM Spectrum Control servers.	✓
Analytics	Performance planning	✓
	Capacity planning	✓
	Business impact analysis (applications, departments, and groups)	✓
	Optimize data placement with tiering	
	Optimize capacity with reclamation.	✓

Restriction: In replication relationships, any snapshot volumes in either source or target storage systems that are not assigned to a host, are not stored in IBM Spectrum Control. The volumes are listed in the Copy Data page but do not have hyperlinks.

Metadata collection schedule

By default, asset, capacity, and configuration metadata is aggregated and collected daily. Schedule daily reports to gain insights about your Hitachi VSP storage systems.

IBM Spectrum Control collects performance metadata from Hitachi VSP storage systems at intervals of 5 minutes and 60 minutes. The default interval is 5 minutes.

Important: When you set the Sample Interval in Hitachi Device Manger to 1 minute or 5 minutes, there will be a gap in the performance data after performance monitoring starts because no data has accumulated. Depending on the value of the Performance Monitor interval in IBM Spectrum Control, the resulting gap in performance data will be 35 minutes or 90 minutes.

What's next

Before you can monitor, alert, and report on your Hitachi VSP storage systems, you must add them to IBM Spectrum Control.

Important: If you want to monitor performance information, you must install the Hitachi Export Tool. The Export Tool requires that the password for connecting to the storage system is temporarily stored as clear text in a file on the server or virtual machine where IBM Spectrum Control is installed. For security reasons, ensure access to that server or virtual machine is restricted to key personnel. The file is automatically deleted after metadata is collected.

For your reference: Want to learn more about support for Hitachi VSP storage systems in IBM Spectrum Control? No problem. Take a look at the following documentation:

- For instructions on how to add storage systems for monitoring, see [Adding storage systems](#).
- For instructions on how to install the Hitachi Export Tool, see [Installing the Hitachi Export Tool](#).
- For instructions on using Hitachi Performance Monitoring and the Hitachi Export Tool, see the Hitachi documentation at <https://knowledge.hitachivantara.com/Documents/Storage>

Related reference

- [Performance metrics for Hitachi VSP storage systems](#)

Planning for NetApp device support

NetApp devices provide unified storage for organizations with network-attached storage or storage area network environments that have file and block-level services. Plan for how to use IBM Spectrum® Control to monitor your NetApp storage systems.

To view the versions of NetApp storage systems that are supported in IBM Spectrum Control, go to the [NetApp support page](#).

Benefits

IBM Spectrum Control can help you predict and prevent storage problems before they impact your business. Here are some key benefits of using IBM Spectrum Control to monitor your NetApp storage systems:

- View detailed information about capacity, storage usage, and performance.
- Monitor health, status, and availability.
- Use alerts and alert policies to be notified of conditions and potential problems.
- Use advanced analytics to reclaim storage.
- Create and share reports about inventory, capacity, performance, and storage consumption.

Device configuration

IBM Spectrum Control can manage different types of network-attached storage. The requirements and device management capabilities are different depending on the type of network-attached storage system you are managing with IBM Spectrum Control.

Device configuration for NetApp storage systems running ONTAP 9

You can add a NetApp storage system running ONTAP 9 to IBM Spectrum Control as either a block or file storage system to view all of the block and file data on the system. When you add a NetApp ONTAP 9 device, IBM Spectrum Control collects data by connecting directly to the storage system.

Antivirus software restriction: If your antivirus software is set on the maximum mode, it might prevent your ONTAP 9 storage systems from being added to IBM Spectrum Control. For more information about how to configure your antivirus software, see [Installation checklists for IBM Spectrum Control](#).

Device configuration for NetApp storage systems running ONTAP 8 (7- mode), or earlier

A NetApp device running ONTAP 8 (7- mode), or earlier, can be added to IBM Spectrum Control as a block storage system, a file storage system, or both.

- When you add a NetApp device running ONTAP 8 (7- mode), or earlier, as a block storage system, IBM Spectrum Control collects data by using an SMI-S agent.
- When you add a NetApp device running ONTAP 8 (7- mode), or earlier, as a file storage system, IBM Spectrum Control collects data through the native API (NAPI) of the device.

Note: For new IBM Spectrum Control installations, Storage Resource agents are no longer supported for managing NetApp devices. If you upgrade from a previous version of IBM Spectrum Control, you can continue to manage NetApp devices that have Storage Resource agents installed from a previous release.

For NetApp devices that you managed with earlier versions of IBM Spectrum Control, you can continue to use the Storage Resource agents that you installed. When you upgrade to IBM Spectrum Control, you can easily upgrade those devices to SMI-S agents (block storage) or network API (file storage). For devices that you continue to manage with Storage Resource agents, use the following guidelines:

- If you configured a Storage Resource agent to be a Scan/Proxy agent, IBM Spectrum Control collects file storage information.
- If you configure a NetApp SMI-S agent (CIM agent), IBM Spectrum Control collects block storage information.
- If you configure a NetApp device as both a filer and a storage system, IBM Spectrum Control collects both file and block storage information.
- If multiple Storage Resource agents are configured to probe the same NetApp filer, the Storage Resource agent that was added to IBM Spectrum Control first is used. Therefore, only data that is gathered by the first Storage Resource agent is shown. For example, if you add a Storage Resource Agent for the Windows operating system and then you add a Storage Resource Agent for the UNIX operating system, Windows operating system-related data such as shares are shown, but not exports.

If you initially configured a NetApp device using a Storage Resource agent in an earlier release of IBM Spectrum Control, you do not lose any function if you decide to also configure the NetApp device by using an SMI-S agent when you upgrade to IBM Spectrum Control.

Supported features

View a detailed list of the features in IBM Spectrum Control that you can use to monitor NetApp ONTAP 9 storage systems:

Table 1. Supported features for NetApp storage systems running ONTAP 9

Resource Monitoring	Features	Supported
Understanding the environment	Monitor storage inventory and configuration. Includes information about type, model, serial number, and firmware.	✓
	Understand storage relationships, from volume and share down to server and application.	✓
	Explore virtualization relationships.	
	Explore replication relationships.	✓
	View dashboards to get insights into key aspects of your storage at a glance and one-click access to web-based element managers.	✓
	Support for multiple protocols and storage types, such as FC, NVMe, CIFS, and NFS.	✓
Monitoring capacity	Collect storage consumption and capacity metrics.	✓
	View data reduction information.	✓
	View copy data information.	✓
	View internal storage tiers such as Easy Tier®.	✓
	Monitor the storage consumed by applications.	✓
Monitoring performance	Collect performance metrics about the workload on resources.	✓
	View calculated metrics to gain insights into performance conditions.	✓
	Export performance data to a compressed file.	✓
	Drill down performance workflows to troubleshoot bottlenecks.	✓
	Compare the performance of resources.	✓
Monitoring health	Understand the health of resources.	✓
	Receive notifications when the status of a resource changes.	✓
	View the status of elements that are not represented as resources.	✓

Resource Monitoring	Features	Supported
Alerting	Alert on conditions within your storage environment.	✓
	Define alerts to identify issues based on multiple conditions.	✓
	Define alert policies to be notified of changes across related resources.	✓
Reporting	View predefined reports.	✓
	View chargeback and consumer reports.	✓
	Create custom reports by using the REST API.	✓
	Create rollup reports to view information across multiple IBM Spectrum Control servers.	✓
Analytics	Performance planning	✓
	Capacity planning	✓
	Business impact analysis (applications, departments, and groups)	✓
	Optimize data placement with tiering	
	Optimize capacity with reclamation.	✓

Restrictions:

- In replication relationships, any snapshot volumes in either source or target storage systems that are not assigned to a host, are not stored in IBM Spectrum Control. The volumes are listed in the Copy Data page but do not have hyperlinks.
- Although you can view the remote and FlashCopy® relationships for NetApp storage systems, the data is not correlated.
- NetApp doesn't provide performance data for any volumes that don't have host connections. Because of this, the number of volumes on the Volumes Performance page might not match the number of volumes on the Volumes page.
- NetApp doesn't provide performance data for any drives that don't have RAID arrays. Because of this, the number of drives on the Drives Performance page might not match the number of drives on the Drives page.

Metadata collection schedule

By default, asset, capacity, and configuration metadata is aggregated and collected daily. You can schedule daily capacity and inventory reports to gain insights about your NetApp storage systems.

For NetApp storage systems running ONTAP 9, the collection intervals for performance metadata are 5 minutes and 60 minutes. The default interval is 5 minutes.

Performance metadata is collected every 15 minutes for other NetApp storage systems.

What's next

Before you can monitor, alert, and report on your NetApp storage systems, you must add them to IBM Spectrum Control.

For instructions on how to add storage systems for monitoring, see [Adding storage systems](#).

Limitations of support for NetApp storage systems running ONTAP 8 (7- mode)

Keep in mind the following limitations for NetApp support:

- The NetApp Data ONTAP SMI-S Agent 3.0 or later implements the Block Server Performance subprofile. It provides volume performance data but for now, does not provide performance data at the storage system level.
- IBM Spectrum Control supports only the SMI-S Array profile of the NetApp Data ONTAP SMI-S Agent. Other SMI-S profiles, including the self-contained NAS profile and the NAS Head profile, are not supported.
- For the Data ONTAP SMI-S 3.0 agent or later, if a volume is offline, the performance monitor might fail with the message: **PM HWNP2132W Performance data could not be collected for device device.**

Related tasks

- [Adding and configuring NetApp resources](#)
- [Removing resources](#)

Related reference

- [Performance metrics for NetApp storage systems](#)

Planning for Pure Storage systems

Pure storage resources provide block storage for organizations with network-attached storage or storage area network environments that have block-level services. IBM Spectrum® Control provides enhanced monitoring for Pure FlashArray//M and FlashArray//X storage systems.

Pure FlashArray//M and FlashArray//X Device configuration

A Pure FlashArray//M or FlashArray//X can be added to IBM Spectrum Control as a block storage system. When you add a Pure storage system, IBM Spectrum Control collects data by connecting directly to the storage system.

To view the versions of Pure storage systems that are supported in IBM Spectrum Control, go to the [Pure Storage support page](#).

Antivirus software restriction: If your antivirus software is set on the maximum mode, it might prevent your Pure Storage systems from being added to IBM Spectrum Control. For more information about how to configure your antivirus software, see [Installation checklists for IBM Spectrum Control](#).

Benefits

IBM Spectrum Control can help you predict and prevent storage problems before they impact your business. Here are some key benefits of using IBM Spectrum Control to monitor your Pure storage systems:

- View detailed information about capacity, storage usage, and performance.
- Monitor health, status, and availability.
- Use alerts and alert policies to be notified of conditions and potential problems.
- Use advanced analytics to reclaim storage.
- Create and share reports about inventory, capacity, performance, and storage consumption.

Supported features

View a detailed list of the features in IBM Spectrum Control that you can use to monitor PureFlashArray//M and FlashArray//X storage systems:

Table 1. Supported features for PureFlashArray//M and FlashArray//X

Resource Monitoring	Features	Supported
Understanding the environment	Monitor storage inventory and configuration. Includes information about type, model, serial number, and firmware.	✓
	Understand storage relationships, from volume and share down to server and application.	✓
	Explore virtualization relationships.	
	Explore replication relationships.	✓
	View dashboards to get insights into key aspects of your storage at a glance and one-click access to web-based element managers.	✓
	Support for multiple protocols and storage types, such as FC, iSCSI, and NVMe.	✓
Monitoring capacity	Collect storage consumption and capacity metrics.	✓
	View data reduction information.	✓
	View copy data information.	✓
	View internal storage tiers such as Easy Tier®.	
	Monitor the storage consumed by applications.	✓
Monitoring performance	Collect performance metrics about the workload on resources.	✓
	View calculated metrics to gain insights into performance conditions.	✓
	Export performance data to a compressed file.	✓
	Drill down performance workflows to troubleshoot bottlenecks.	✓
	Compare the performance of resources.	✓
Monitoring health	Understand the health of resources.	✓
	Receive notifications when the status of a resource changes.	
	View the status of elements that are not represented as resources.	✓
Alerting	Alert on conditions within your storage environment.	✓
	Define alerts to identify issues based on multiple conditions.	✓
	Define alert policies to be notified of changes across related resources.	✓
Reporting	View predefined reports.	✓
	View chargeback and consumer reports.	✓
	Create custom reports by using the REST API.	✓
	Create rollup reports to view information across multiple IBM Spectrum Control servers.	✓
Analytics	Performance planning.	✓
	Capacity planning.	✓
	Business impact analysis (applications, departments, and groups).	✓
	Optimize data placement with tiering.	
	Optimize capacity with reclamation.	✓

Restrictions:

- In replication relationships, any snapshot volumes in either source or target storage systems that are not assigned to a host, are not stored in IBM Spectrum Control. The volumes are listed in the Copy Data page but do not have hyperlinks.
- For asynchronous remote relationships, note the following limitations.
 - From the target storage system perspective, no details are available for the source storage system.
 - From the source storage system perspective, the remote storage systems and volumes are listed in the Copy Data page. However, the remote volumes cannot be correlated with the volumes on the source system and so do not have hyperlinks.

Metadata collection schedule

By default, asset, capacity, and configuration metadata is aggregated and collected daily. Schedule daily reports to gain insights about your Pure storage systems.

For Pure storage systems, the collection intervals for performance metadata are 5 minutes and 60 minutes. The default interval is 5 minutes.

What's next

Before you can monitor, alert, and report on your Pure storage systems, you must add them to IBM Spectrum Control.

For instructions on how to add storage systems for monitoring, see [Adding storage systems](#).

Related reference

- [Performance metrics for Pure storage systems](#)

Planning to monitor performance

IBM Spectrum® Control can collect performance metrics for resources that use the native interfaces, such as DS8000®, XIV®, SAN Volume Controller, Storwize® V7000 Unified, or Storwize V7000 storage systems, or for storage systems and Fibre Channel switches that use SMI-S providers (also called CIM agents or CIMOMs). The resources that use SMI-S providers must be SMI-S 1.1 compliant. To collect performance metrics for IBM FlashSystem® 900, the SNMP agent must be enabled on the storage system.

About this task

You can monitor the performance of storage resources by completing the following tasks:

- Schedule performance monitors to collect performance information about a resource.
- Use the performance alerts in the default alert policies or create alerts that notify when performance thresholds are met or exceeded for a resource.
- View performance metrics, charts, and reports about a resource.

You can use a combination of these tasks to achieve the following goals:

- Monitor a complicated storage network environment.
- Predict warning signs of system failure.
- Do capacity planning as overall workload grows.
- [Performance metrics](#)
IBM Spectrum Control can collect many different performance metrics, which indicate the particular performance characteristics of monitored resources.
- [Performance configuration](#)
Performance data can be collected from resources that are managed through the native interfaces or from resources that are managed by an SMI-S provider (also called CIM agent or CIMOM). Resources that use an SMI-S provider must be SMI-S compliant.

Performance metrics

IBM Spectrum® Control can collect many different performance metrics, which indicate the particular performance characteristics of monitored resources.

About this task

Two important metrics for storage systems are throughput in I/O per second, and the response time in milliseconds. Throughput is measured and reported in several different ways:

- Throughput of an entire box (storage system)
- Each cluster
- Each controller (Examples: DS8000®)
- Each I/O Group (Example: storage systems that run IBM Spectrum Virtualize)

Throughputs are measured for:

- Each volume (or LUN)
- At the Fibre Channel interfaces (ports) on some of the storage boxes
- On Fibre Channel switches
- At the RAID array after cache hits have been filtered out

For storage systems, the performance statistics are separated into frontend I/O metrics and back-end I/O metrics. Front-end I/O metrics are a measure of the traffic between the servers and storage systems. Back-end I/O metrics are a measure of all traffic between the storage system cache and the disks in the RAID arrays in the back-end of the storage system. Most storage systems give metrics for both kinds of I/O operations: front-end and back-end operations. It is important to know whether the throughput and response times are at the front-end (close to the system level response time as measured from a server) or back-end (between the cache and disk).

The main front-end throughput metrics are:

- Total IO rate (overall)
- Read IO rate (overall)
- Write IO rate (overall)

The corresponding front-end response time metrics are:

- Overall response time
- Read response time
- Write response time

The main back-end throughput metrics are:

- Total back-end IO rate (overall)
- Back-end read IO rate (overall)
- Back-end write IO rate (overall)

The corresponding back-end response time metrics are:

- Overall back-end response time
- Back-end read response time
- Back-end write response time

For planning purposes, it's important to track any growth or change in the rates and response times. It frequently happens that I/O rate grows over time, and that response time increases as the I/O rates increase. This relationship is what "capacity planning" is all about. As I/O rates increase, and as response times increase, you can use these trends to project when additional storage performance (as well as capacity) is required.

Depending on the particular storage environment, it might be that throughput or response time times change drastically from hour to hour or day to day. There might be periods when the values fall outside the expected range of values. In that case, other performance metrics can be used to understand what is happening. Here are some additional metrics that can be used to make sense of throughput and response times:

- Total cache hit percentage
- Read cache hit percentage
- Write-cache delay percentage (previously known as NVS full percentage)
- Read transfer size (KB/operation)
- Write transfer size (KB/operation)

Low cache hit percentages can drive up response times, because a cache miss requires access to back-end storage. Low hit percentages also tend to increase the utilization percentage of the back-end storage, which might adversely affect the back-end throughput and response times. High write-cache delay percentage (previously known as NVS full percentage) can drive up the write response times. High transfer sizes typically indicate more of a batch workload, in which case the overall data rates are more important than the I/O rates and the response times.

All these metrics can be monitored through lists, charts, and reports in IBM Spectrum Control. Some examples of supported thresholds are:

- Total I/O rate and total data rate thresholds
- Total back-end I/O rate and total back-end data rate thresholds
- Read back-end response time and write back-end response time thresholds
- Total port I/O rate (packet rate) and data rate thresholds
- Overall port response time threshold
- Port send utilization percentage and port receive utilization percentage thresholds
- Port send bandwidth percentage and port receive bandwidth percentage thresholds

For Fibre Channel switches, the important metrics are total port packet rate and total port data rate, which provide the traffic pattern over a particular switch port. Port bandwidth percentage metrics are also important to provide an indicator of bandwidth usage based on port speeds. When there are lost frames from the host to the switch port, or from the switch port to a storage device, the dumped frame rate on the port can be monitored.

The important things are:

- Monitor the throughput and response time patterns over time for your environment
- Develop an understanding of expected behaviors
- Investigate the deviations from normal patterns of behavior to get warning signs of abnormal behavior
- Generate the trend of workload changes

Related reference

- [Performance metrics](#)

Performance configuration

Performance data can be collected from resources that are managed through the native interfaces or from resources that are managed by an SMI-S provider (also called CIM agent or CIMOM). Resources that use an SMI-S provider must be SMI-S compliant.

Procedure

To monitor the performance of resources, you must complete the following tasks:

1. Prepare for performance monitoring.
For resources that require an SMI-S provider, ensure that the following conditions are met before you add the SMI-S provider to IBM Spectrum® Control:
 - The version of the SMI-S provider and the firmware for the device is supported. For information about the SMI-S provider and firmware that is supported, see <https://www.ibm.com/support/pages/node/388393> and go to the *Storage* section.
 - An SMI-S provider is installed on a different server than the IBM Spectrum Control server. For some switch vendors, the switch comes with an embedded SMI-S provider, so an SMI-S provider does not need to be installed.
 - For storage systems or switches on a private network, ensure that the SMI-S provider is installed on a gateway machine so that the IBM Spectrum Control server can communicate with that agent.
 - The SMI-S provider is configured to manage the intended resource.
2. Add resources for monitoring and schedule data collection.
Before you can view the performance of a resource, you must add it to IBM Spectrum Control for monitoring. When you add resources, you can automatically schedule performance monitors. Performance monitors are data collection jobs that collect performance information.
For more information about adding resources, see [Adding resources](#) and search for *adding resources*.
3. Use the performance alerts in the default alert policies or create alerts for performance thresholds.
To be notified when the performance of a resource might represent a potential problem, define a threshold alert to monitor a specific metric. In a threshold alert, you can specify one or more boundary values for a metric. When the performance of a resource falls outside these boundaries, an alert is generated. For example, you can define an alert threshold that notifies you when the total I/O rate for a port on a storage system falls outside a specified range. Chronologically, you must define a threshold alert before a performance monitor starts collecting data.
Alerts are triggered by conditions and violations of performance thresholds that are detected during data collection and event processing. For some storage systems such as IBM Spectrum Accelerate and the XIV®, events are polled every minute from the resource. For IBM Spectrum Scale, status change events are

polled frequently, typically within minutes. For other resources, events are subscription-based, where the resource itself or a data source such as an SMI-S provider sends the events to IBM Spectrum Control when conditions change on the resource. Examples of storage systems that use subscription-based event processing include SAN Volume Controller, Storwize® V7000, Storwize V7000 Unified, and FlashSystem V9000.

Tip: Alerts can be created, deleted, or modified at any time, even if the performance monitor is already running. Changes are applied dynamically to the running performance monitor and take effect the next time that data is collected.

4. Check the status of a performance monitor.

When a performance monitor starts to run, IBM Spectrum Control begins to collect performance data for the specified resource. To check the status of the performance monitor to ensure that it is running or completed successfully, complete the following steps:

- In the menu bar, go to Home > Performance Monitors.
- In the Name column, locate the name of the resource that was included in the performance monitor.
- In the Status column, check the status of the performance monitor.

5. Specify how long performance data is retained.

In the GUI, go to Settings > History Retention to configure how long to retain performance data that IBM Spectrum Control collects about resources. You can specify a retention period for collected sample data, for aggregated hourly data, and for daily data. The retention period applies to all performance monitors. The longer that you keep data, the more informative your analysis.

Sample data is the data that is collected at the specified interval length of the performance monitor. For example, data is collected every 5 minutes. For storage systems, the most numerous components are usually volumes, therefore the largest amount of performance data is collected for volumes.

- [Adding resources](#)
- [Collecting data](#)
- [Alerting](#)
- [Monitoring the performance of resources](#)
- [Configuring history and data retention](#)

Planning for switches and fabrics

IBM Spectrum® Control supports IBM®, Brocade, and Cisco switches. Use IBM Spectrum Control to help manage the switches and fabrics that connect host systems and applications to storage resources.

Switches

IBM Spectrum Control can monitor the following switches:

Brocade

IBM Spectrum Control can monitor Brocade switches and fabrics in either of the following ways:

- For switches with Fabric OS 8.2.1 or later, connect directly to the switches and use the REST API to manage the switches and fabrics.
- For switches with a Fabric OS version earlier than 8.2.1, use the Brocade Network Advisor (BNA) Storage Management Initiative (SMI) agent to manage the switches and fabrics. The SMI agent is embedded in the BNA. BNA can manage multiple fabrics within and across data centers. When you configure BNA, you set up one switch to be the *seed* switch that interconnects to all the other switches in the fabric. Important: The embedded SMI agent is only available in the Professional Plus and Enterprise editions of BNA. BNA is no longer available for purchase, although it will be supported by Broadcom until February 2022. For more information about end of support for BNA, see <https://www.broadcom.com/support/fibre-channel-networking/eol>.

Cisco

IBM Spectrum Control uses Simple Network Management Protocol (SNMP) agents to discover and collect data about Cisco switches and fabrics in your environment.

For the most current information about switches and directors supported, see the IBM Spectrum Control support website at <https://www.ibm.com/support/pages/node/6249365>.

Fabrics

The main IBM Spectrum Control functions for managing fabrics are shown in the following table.

Table 1. Main fabric management features

Feature	Advantages	Benefit
Monitoring and alerts	Monitor SAN events and alerts administrators of problems	Helps maintain the availability of your SAN
SAN reporting	Generate reports on the SAN devices	Quickly provides reports and inventories of your SAN
Enterprise scalability	Scale from SAN islands to enterprise SANs	Accommodates new and ever changing business needs
Monitoring performance	Monitor the performance of switches by defining performance alerts and viewing performance reports. The performance reporting is at the switch and port levels.	Helps you maintain high SAN availability

Note:

1. IBM Spectrum Control provides predictive fault analytics for the SAN infrastructure through reporting, monitoring, and alerting of link failure rates and error frame rates. These transient errors are indicative of potential pending SAN link failures.
2. IBM Spectrum Control 5.1 or later do not support heterogeneous fabrics.

Monitoring fabrics and switches

Switches must be added to IBM Spectrum Control before they can be monitored. When you add the switch for monitoring, ensure that you schedule a probe and performance monitor to collect configuration and performance information about that switch.

For information about adding switches and fabrics for monitoring, see [Adding fabrics and switches](#).

- **[Collecting data about fabrics](#)**
IBM Spectrum Control uses REST APIs, SMI agents, and SNMP agents to collect data about Brocade and Cisco fabrics and switches.
- **[Planning for Brocade](#)**
IBM Spectrum Control can monitor Brocade switches and fabrics by connecting directly to the switches, or by using the SMI agent that is embedded in Brocade Network Advisor (BNA).
- **[Planning for Cisco](#)**
IBM Spectrum Control uses Simple Network Management Protocol (SNMP) data sources to manage Cisco switches and fabrics.
- **[Migrating fabrics and switches](#)**
IBM Spectrum Control can monitor Brocade or Broadcom switches and fabrics by connecting directly to the switches, or by using the SMI agent that is embedded in Brocade Network Advisor (BNA). IBM Spectrum Control can also monitor Cisco switches and fabrics by using an SNMP agent.
- **[Planning for private switch networks](#)**
Some switch vendors recommend a private IP network for Fibre Channel (FC) switches.

Collecting data about fabrics

IBM Spectrum® Control uses REST APIs, SMI agents, and SNMP agents to collect data about Brocade and Cisco fabrics and switches.

About this task

The actions that trigger data collection, and the data that is collected, depend on how the switch is monitored.

Brocade switches with Fabric OS 8.2.1 or later: REST API

You connect directly to Brocade switches that are running Fabric OS 8.2.1 or later, and monitor the switches by using the Fabric OS REST API. Data is collected in the following ways:

- Probe jobs collect asset and status information about the switches. You can schedule probes or start them when you need them to collect data immediately. To collect complete asset and status information about a fabric, every switch on the fabric must be probed.
- Performance monitors collect metrics about the performance of switches. You can schedule performance monitors for switches only, not for fabrics.

IBM Spectrum Control does not collect event-driven data for Brocade switches that are running Fabric OS 8.2.1 or later.

Brocade switches with a Fabric OS version earlier than 8.2.1: Network Advisor SMI agent

Brocade switches use the Brocade Network Advisor (BNA) SMI agents to send CIM indications to the IBM Spectrum Control server. IBM Spectrum Control automatically subscribes to these CIM indications, so you do not need to do any additional configuration to receive these notifications.

The SMI agent does a mini-probe to collect information that is relevant to the indication received. The SMI agent might do one of the following actions:

- Retrieve switches and fabrics.
- Retrieve zoning data.
- Retrieve switch port status and connection to node.
- Retrieve switch blade status and all port connections for the blade.
- Set fabric, switch, blades, connections, or nodes as missing.

Cisco switches: SNMP agent

An SNMP trap is received from a Cisco switch, and this triggers data collection. The SNMP agent performs a probe of fabrics for which the traps were received. The SNMP agent does the following actions:

- Retrieves switch and topology information. To collect complete asset and status information about a fabric, every switch on the fabric must be probed.
- Retrieves Cisco VSAN and physical infrastructure and correlates VSAN to physical infrastructure.
- Gathers zoning information and provides zone control support for Cisco switches.

For information about the supported agent types by switch vendor, see [Planning for Brocade](#) or [Planning for Cisco](#).

Planning for Brocade

IBM Spectrum® Control can monitor Brocade switches and fabrics by connecting directly to the switches, or by using the SMI agent that is embedded in Brocade Network Advisor (BNA).

Supported Brocade models

To see the Brocade models, firmware versions, and agents that are used by IBM Spectrum Control, see the [IBM Spectrum Control interoperability matrix](#).

Monitoring fabrics and switches

IBM Spectrum Control can monitor Brocade switches and fabrics in either of the following ways:

- For switches with Fabric OS 8.2.1 or later, connect directly to the switch and use the REST API to manage the switches and fabrics.
IBM Spectrum Control connects to the switch that you added, and automatically discovers the fabrics and switches that they manage. The discovered fabrics and switches are automatically added for monitoring, if they have the same credentials as the switches that you added. When you add switches, you can schedule probes and performance monitors to collect data for your switches and fabrics.
- For switches with a Fabric OS version earlier than 8.2.1, use the BNA Storage Management Initiative (SMI) agent to manage the switches and fabrics. The SMI agent is embedded in the BNA. BNA can manage multiple fabrics within and across data centers. When you configure BNA, you set up one switch to be the *seed* switch that interconnects to all the other switches in the fabric.

When you add fabrics and switches for monitoring, you can specify the SMI agent that manages them. IBM Spectrum Control connects to the agent and automatically discovers the fabrics and switches that it manages.

BNA can manage multiple fabrics within and across data centers. When you configure BNA, you set up one switch to be the *seed* switch that interconnects to all the other switches in the fabric.

Important: The embedded SMI agent is only available in the Professional Plus and Enterprise editions of BNA. BNA is no longer available for purchase, although it is supported by Broadcom until February 2022. For more information about end of support for BNA, see <https://www.broadcom.com/support/fibre-channel-networking/eol>.

Updating BNA: You can access Brocade FOS firmware and Network Advisor *updates* (not new installations) through IBM® Fix Central. You no longer need to open a case with IBM support to request them. For more information, see [Brocade FOS and Network Advisor updates](#).

Choosing whether to use REST API or BNA to monitor switches

Use the following table to help decide how to monitor your Brocade switches:

Monitoring method	Use in these situations
Direct REST API (No other software is required)	<ul style="list-style-type: none">• BNA is not available to you.• You are not monitoring any switches or fabrics yet.• Your models of Brocade switches are not supported by BNA.• You are adding a new fabric to IBM Spectrum Control. <p>Consider the following additional information when you are connecting directly to switches:</p> <ul style="list-style-type: none">• If you are connecting directly to Brocade switches, you need to add only one switch from each fabric to IBM Spectrum Control. If other switches that are on the same fabric are running Fabric OS 8.2.1 or later and have the same username, password, protocol, and port, they are added to IBM Spectrum Control automatically. The switches are probed automatically and the performance monitoring starts running.• Switches that have different authentication credentials or that are running a version of Fabric OS earlier than 8.2.1, are added to the Switches page with a Not Monitored condition. You can add the switches that are not monitored by following the usual process to add switches.• If you are monitoring a switch through BNA and want to use the REST API instead, you must first remove that switch and then add it again for monitoring. When you remove a switch, its history of performance metadata is automatically deleted and no longer available.
SMI-S (Network Advisor is required)	<ul style="list-style-type: none">• You are already monitoring switches by using BNA.• You are adding switches that are supported by BNA.

Tip:

- After you add a switch, you can modify its connection. For example, if you add a switch by using its IP address, or a switch was automatically discovered and added by IBM Spectrum Control using an IP address, you can later modify its connection to use a host name instead. To save time, if you want to collect metadata for multiple switches in the same way (host name or IP address), it is recommended that you use that method when you add the switches for monitoring.
- When you add switches through the REST API, the Brocade model name for IBM-branded OEM switch models is shown in the GUI. To help you identify a switch with IBM support, edit the Custom Tag in the properties for that switch.

Restriction:

- BNA is supported by Broadcom until February 2022. However, support for new models of switch will not be added to the SMI agent.
- You cannot manage a mixture of directly connected switches and switches that are managed by BNA on the same fabric. It is recommended that you move switches that are monitored through BNA to a separate fabric from REST-enabled switches.
- If you are adding switches that are not supported by BNA, then you must connect to the new switches directly and use Fabric OS REST API to manage them.

Restrictions for connecting directly to Brocade switches

Consider the following information when you are planning to monitor directly connected switches that are running Fabric OS 8.2.1 or later:

- Brocade switches and their associated fabrics can be monitored only through one data source at a time. If you upgrade from an earlier version of IBM Spectrum Control, you must remove the fabrics that are monitored through BNA before you can add the switches in those fabrics by direct connection.
- To schedule probes for Brocade switches, use the Switches page, not the Fabrics page.
- ISLs and trunks are shown for switches only after both the switch and the connected switch are probed.
- In rollup mode, information is only displayed for the default logical switch when virtualization is enabled on the switch. Information is not displayed for the physical switch.
- You can't receive asynchronous events from your switches if you are monitoring your fabrics and switches by using Fabric OS REST API.
- [Planning to use the BNA SMI agent](#)
The SMI agent, sometimes known as a CIM agent, is embedded in the Brocade Network Advisor (BNA). BNA manages multiple fabrics within and across data centers. When you configure BNA, you set up one switch to be the *seed* switch that interconnects to all the other switches in the fabric.
- [Support for Access Gateway switches and virtual fabrics](#)
IBM Spectrum Control can monitor Brocade Access Gateway switches and virtual fabrics.

Related concepts

- [Agents](#)

Related reference

- [Data sources for switches and fabrics](#)

Planning to use the BNA SMI agent

The SMI agent, sometimes known as a CIM agent, is embedded in the Brocade Network Advisor (BNA). BNA manages multiple fabrics within and across data centers. When you configure BNA, you set up one switch to be the *seed* switch that interconnects to all the other switches in the fabric.

Configuring the Brocade switch

Before IBM Spectrum® Control can discover the switches in a fabric through the SMI agent, you must configure BNA to manage the fabric. Use a seed switch to configure BNA to manage the fabric.

There might be significant limitations on managing fabrics, depending on the version of BNA that you are running. Check the [IBM Spectrum Control interoperability matrix for switches](#) for information about the switches and directors that are supported by IBM Spectrum Control, and limitations that you need to know about when you use BNA.

Important: The embedded SMI agent is only available in the Professional Plus and Enterprise editions of BNA. BNA is no longer available for purchase, although it will be supported by Broadcom until February 2022. For more information about end of support for BNA, see <https://www.broadcom.com/support/fibre-channel-networking/eol>. Updating BNA: You can access Brocade FOS firmware and Network Advisor *updates* (not new installations) through IBM® Fix Central. You no longer need to open a case with IBM support to request them. For more information, see [Brocade FOS and Network Advisor updates](#).

Support for Access Gateway switches and virtual fabrics

IBM Spectrum® Control can monitor Brocade Access Gateway switches and virtual fabrics.

Support for Brocade Access Gateway switches

IBM Spectrum Control discovers and monitors Brocade Access Gateway switches, also called N_Port Virtualizer (NPV) switches. These switches are not probed as part of the fabric. Each Brocade Access Gateway switch has a separate probe job defined. The probe collects the following information:

- The switch properties.
- The ports of the switch.
- The connections to these switch ports.

Support for Brocade virtual fabrics

You can monitor Brocade virtual fabrics, with some limitations. To monitor Brocade virtual fabrics, you can:

- Discover Brocade virtual fabrics and switches.
- Schedule probes to collect asset and status data about Brocade virtual fabrics.
- Schedule monitors to collect information about the performance of Brocade virtual switches.

When you add and collect information about Brocade virtual fabrics, you can:

- See the properties of Brocade virtual fabrics and switches and the resources that are connected to them. The virtual fabrics and switches are displayed as if they are physical fabrics and switches.
- Control zones for virtual fabrics. Zones are created during storage provisioning.

The following limitations apply to the support for Brocade virtual fabrics that use Brocade Network Advisor (BNA):

- If Brocade switches are partitioned into virtual fabrics, the virtual fabrics and switches are displayed in the IBM Spectrum Control GUI, but the physical fabrics and physical switches are not displayed.
- The virtual fabrics and switches are displayed as if they are physical fabrics and switches. Therefore, you cannot see which virtual fabrics are on the same physical fabric, and you cannot see which virtual switches are on the same physical switch.
- In the IBM Spectrum Control GUI and in Cognos® Analytics reports, the Brocade virtual fabrics and switches are incorrectly identified as physical fabrics and switches.

There might be significant limitations on managing fabrics, depending on the version of BNA that you are running. Check the [IBM Spectrum Control interoperability matrix for switches](#) for information about the switches that are supported by IBM Spectrum Control, and any limitations that you need to know about when you use BNA.

Planning for Cisco

IBM Spectrum® Control uses Simple Network Management Protocol (SNMP) data sources to manage Cisco switches and fabrics.

Supported Cisco models

IBM Spectrum Control supports a number of different Cisco switch models and firmware versions. For information about the supported versions, see the **IBM Spectrum Control 5.3.x - Switches - Supported Products Matrix** at <https://www.ibm.com/support/pages/node/6249365>.

Consider the following information when you are using Cisco switches with IBM Spectrum Control:

- The Cisco Nexus 5000 switch is supported in toleration mode only. IBM Spectrum Control does not support the Converged Enhanced Ethernet (CEE) or Fibre Channel over Ethernet (FCoE) connectivity functions for the Cisco Nexus 5000 switch.

- The Fibre Channel (FC) or Fibre Channel over Ethernet (FCoE) protocols must be enabled on the switch. Some switches, such as the Cisco Nexus 5000 series, require you to enable these protocols. Otherwise, IBM Spectrum Control will not recognize the switch when you try to add it using the Add Switches and Fabrics for Monitoring dialog. For instructions on how to configure Cisco switches for FCoE enablement, go to the Cisco product website at <http://www.cisco.com> and click Support.
- IBM Spectrum Control also does not support domain-port zoning for Cisco Nexus switches. IBM Spectrum Control displays the FCoE and FC ports in the switch port lists.
- The password that IBM Spectrum Control uses to connect to a Cisco switch can't contain these special characters: < >.

Supported fabric management functions

The following fabric management functions are supported by SNMP data sources, also referred to as SNMP agents:

- Collecting information about Cisco switches and switch ports.
- Collecting information about topology connectivity.
- Collecting information about zoning information and zone control.
- Monitoring performance.
- Generating alerts.
- **Cisco SNMP agent**
IBM Spectrum Control uses SNMP queries to discover information about the SAN. Management Information Base (MIB) information is collected from the switches and directors by the SNMP agent. Switches and directors can be added as SNMP agents and contacted from the IBM Spectrum Control Device server by SNMP.
- **Configuration requirements for Cisco switches and directors**
IBM Spectrum Control uses SNMP agents to collect data about Cisco switches and fabrics in the SAN.

Related concepts

- [Agents](#)

Cisco SNMP agent

IBM Spectrum® Control uses SNMP queries to discover information about the SAN. Management Information Base (MIB) information is collected from the switches and directors by the SNMP agent. Switches and directors can be added as SNMP agents and contacted from the IBM Spectrum Control Device server by SNMP.

SNMP agent functions

The SNMP agent performs the following functions:

- Gathers information about the fabric by querying the switch or director.
- Gathers virtual SAN information for Cisco switches.
- Gathers zoning information and provides zone control support for Cisco switches.

Important note: Topology information is only gathered for the switch added as an SNMP agent. The agent cannot gather the topology information for any connected switches unless they are also defined as SNMP agents. If you are performing discovery on a fabric with several switches, you must install an SNMP agent for each switch in the fabric to discover the whole fabric.

Limitations:

- Information about fabric node devices such as servers and storage devices is limited. Most devices are unknown with a type of "Other" and are identified by their WWN.
- There must be a TCP/IP connection between the switch and the IBM Spectrum Control Device server.
- To enable events from the switch to the Device server, the switch must be configured to send SNMP traps to the Device server.

Determining why an SNMP agent shows a state of "not contacted"

Complete the following tasks to determine why an SNMP agent shows a state of "not contacted":

- Check that the IP address or host name is entered correctly in the switch properties page.
- Ensure that the Fibre Alliance FC Management MIB is enabled on the switch. IBM Spectrum Control uses the SNMP protocol to send queries across the IP network to management information bases (MIBs) supported on the switch. IBM Spectrum Control uses the Fibre Alliance FC Management MIB and the Fibre Channel FE MIB.
- Ensure that the server IP address is included in any SNMP Access Control List that exists on the switch.
- The switch must be configured to receive SNMP queries and respond using SNMP. SNMPv3 is the preferred version. Some switches are configured to use earlier versions of SNMP by default, and should be configured to use SNMPv3 if possible.

Configuration requirements for Cisco switches and directors

IBM Spectrum® Control uses SNMP agents to collect data about Cisco switches and fabrics in the SAN.

For IBM Spectrum Control to gather and display information from Cisco switches and directors, the SNMP communication protocol must be configured correctly.

Incorrect switch configuration can lead to missing information and misconceptions about the IBM Spectrum Control product not working properly with certain switches.

For a Cisco switch to successfully receive and respond to queries from IBM Spectrum Control, the following basic requirements must be met:

- IBM Spectrum Control can use SNMPv3 (preferred) or SNMPv1 to probe switches and fabrics. The SNMPv3 protocol is preferred because it provides better security, but switches that use the SNMPv1 protocol are also supported. Some switches are configured to use SNMPv3 by default.
- The Fibre Alliance FC Management MIB (FA MIB) and Fibre Channel Fabric Element MIB (FE MIB) must be enabled on the switch.
- When you use the SNMPv1 protocol, the community string that is configured in IBM Spectrum Control must match one of the community strings that are configured on the switch with read access. Additionally, Cisco switches must have a community string match for write access. The default community strings in IBM Spectrum Control are "public" for read access and "private" for write access. Other community strings can be defined on the switches, but are not used. Community strings are not relevant when you use the SNMPv3 protocol.
- SNMP access control lists must include the IBM Spectrum Control system. These access control lists are defined and configured on the switches. Some lists automatically include all hosts, while others exclude all by default.
- The Fibre Channel (FC) or Fibre Channel over Ethernet (FCoE) protocols must be enabled on the switch. Some switches, such as the Cisco Nexus 5000 series, require you to enable these protocols. Otherwise, IBM Spectrum Control does not recognize the switch when you try to add it using the Add Switches and Fabrics for Monitoring dialog. For instructions on how to configure Cisco switches for FCoE enablement, go to the Cisco product website at <http://www.cisco.com> and click Support.

You can configure Cisco switches to notify IBM Spectrum Control of changes in the SAN that occur between your scheduled data collection probes. These notifications, called asynchronous event notifications, are processed as they are received by IBM Spectrum Control. If you want the Cisco switches to send these event notifications to IBM Spectrum Control, configure one or more switches to send SNMP traps to the server.

SNMP traps are generated by Cisco switches and directed to the IBM Spectrum Control Device server. The SNMP trap is an indication that something changed in the fabric, and that a discovery must occur to identify the changes. The default configuration for handling switch traps is to send them from the switch to port 162 on the IBM Spectrum Control Device server. For the successful generation and reception of SNMP traps, the following configuration is required:

- The trap destination must be specified on the switch. The trap destination is typically the IBM Spectrum Control Device server, but it can also be an intermediary SNMP manager that receives the trap and sends it to IBM Spectrum Control.
- The destination port must be specified on the switch. IBM Spectrum Control listens on port 162 by default.
- The traps must be sent as SNMPv3 or SNMPv1. The SNMP designation is set on the switch.
- The trap severity level must be set to generate traps for change conditions. The trap severity level indicates that error level traps and anything more severe are sent. The trap severity level is set on the switch.

Migrating fabrics and switches

IBM Spectrum® Control can monitor Brocade or Broadcom switches and fabrics by connecting directly to the switches, or by using the SMI agent that is embedded in Brocade Network Advisor (BNA). IBM Spectrum Control can also monitor Cisco switches and fabrics by using an SNMP agent.

You can migrate Brocade switches that use the Brocade Network Advisor (BNA) SMI agent to use the Fabric OS REST API, if the switch is running Fabric OS 8.2.1 or later.

Important: The embedded SMI agent is only available in the Professional Plus and Enterprise editions of BNA. BNA is no longer available for purchase, although it is supported by Broadcom until February 2022. For more information about end of support for BNA, see <https://www.broadcom.com/support/fibre-channel-networking/eol>.

Scenario	Considerations
Brocade switches with a version of Fabric OS earlier than 8.2.1 that are managed by using BNA SMI agent.	<p>To manage the switches by using the Fabric OS REST API, you must update the switches to Fabric OS 8.2.1 or later. If you want to use HTTPS to communicate with the switch, you must enable HTTPS on the switch. For more information about how to configure the switch, see the documentation at:</p> <ul style="list-style-type: none"> • IBM® Storage Networking Fibre Channel SAN b-type family documentation • IBM Storage Networking SAN96B-5 Product documents • Broadcom or Brocade Fibre Channel Networking Switches
Brocade switches with Fabric OS 8.2.1 or later that are managed by using BNA SMI agent.	<p>The BNA SMI agent is supported by Broadcom until February 2022, so you can continue to use BNA to manage your switches and fabrics. Use BNA if you want to retain the following benefits:</p> <ul style="list-style-type: none"> • Continue to receive asynchronous events from the switches. • Retain historical performance and error metric data for the switches. • Retain IBM Spectrum Control customizations for the switches, such as reports, alerts, custom properties, and group membership. <p>Your organization might upgrade your switches before the end of support of BNA. In this case, consider continuing to use BNA. If your organization gets new switches, connect to the switches directly and monitor the switch by using the Fabric OS REST API.</p> <p>Note: You cannot manage a mixture of switches that are directly connected and switches that are managed by BNA on the same fabric.</p> <p>If you must add switches that are not supported by BNA, then you must connect to the new switches directly and use the Fabric OS REST API to manage them. Therefore, you must remove the switches from IBM Spectrum Control. Then, add the switches to IBM Spectrum Control again. Select the 8.2.1 or later value for the Fabric OS version.</p> <p>Restriction: If you remove a switch, the asset and performance data is also removed from IBM Spectrum Control. If you defined alerts for a switch that is removed, the alerts are removed also. However, if the alert configuration for the switch was defined in an alert policy, you can apply the alert policy after you connect directly to the switch.</p>

Migrating switches and fabrics from using BNA to connecting directly with REST API

To migrate Brocade switches that use the BNA SMI agent to connect directly to the switches and use the Fabric OS REST API, follow these steps:

1. Remove the switches and fabrics from IBM Spectrum Control.
Restriction: If you remove a switch, the asset and performance data is also removed from IBM Spectrum Control.
2. Remove the switches and fabrics from BNA.
3. Update the version of Fabric OS to 8.2.1 or later on the switches.
4. Add the switches to IBM Spectrum Control.

Tip: You need to add only one switch from each fabric to IBM Spectrum Control. If other switches that are on the same fabric are running Fabric OS 8.2.1 or later and have the same username, password, protocol, and port, they are added to IBM Spectrum Control automatically.

Related tasks

- [Adding fabrics and switches](#)
- [Removing resources](#)

Planning for private switch networks

Some switch vendors recommend a private IP network for Fibre Channel (FC) switches.

If you are planning a private IP network for FC switches, consider the following information:

- The IBM Spectrum® Control server can't communicate with the switches if they are on a private IP network.
- SNMP agents for Cisco fabrics require a TCP/IP connection from the IBM Spectrum Control server to the switch.
- SNMP traps from Cisco switches can't travel directly from the switches to the IBM Spectrum Control server.

To collect SAN data and forward SAN events to the IBM Spectrum Control server, you must enable the server to communicate with the private switch network. Use a second network interface card (NIC) to enable the server to communicate with the private switch network.

Planning for VMware

IBM Spectrum® Control supports the VMware vSphere components ESXi and vCenter Server.

Overview

VMware ESXi is a true hypervisor product that can host multiple virtual machines that run independently of each other while sharing hardware resources. The vCenter Server is the management application that is the central entry point for the management and monitoring of the ESXi hosts for a data center. A Storage Resource agent must be installed on each virtual machine that is hosted by ESXi.

For more information about ESXi or vCenter Server, see <http://www.vmware.com>.

For information about permissions required by IBM Spectrum Control to access information from VMware, see [Checking permissions to browse data stores](#).

For information about the storage systems and platforms supported on VMware, see [Hypervisors and VMware® monitoring support](#).

For full function, both the Storage Resource agent and the ESXi host must be up and running. If one of the items is not present in a given environment, only a limited picture is presented to the user. Some virtual machines might not be recognized.

The hierarchical mapping of storage allocated to the virtual machine is available for the virtual machines on the ESXi host.

Software requirements

For information about the software requirements for your VMware ESXi or vCenter Server environment, see [Software requirements for operating systems](#).

The IBM Spectrum Control server on a virtual machine

You can install the IBM Spectrum Control server on a virtual machine on VMware ESX server 3.5.x. The hardware and operating system requirements are the same requirements as for a physical machine.

In addition, the following requirements must be met:

CPU

For the ESX server, do not have more virtual processors than there are physical cores in the system. Plan your system so that no processor scheduling is required by the VM kernel for the virtual machine.

RAM

Ensure that you have enough RAM in the ESX server to service all the virtual machines with a maximum RAM usage. Plan your system so that the ESX server does not need to swap RAM for the virtual machine.

Disk

Use the SAN-attached RDM with SCSI pass-through for DB2® data and log storage.

Important: To prevent high virtual disk latency and poor virtual machine performance, you must use a dedicated VMDK (Virtual Machine Disk) for the virtual machine where you plan to install IBM Spectrum Control. Do not use a shared VMDK for the virtual machine. When you use a shared VMDK, it reduces the total number of disk IOPS (Input/Output Operations Per Second) available to the virtual machine.

For more information, see <https://kb.vmware.com/s/article/1031773>.

For more information about VMware, see the following publications:

- "Using VMware ESX Server with IBM® WebSphere® Application Server" at http://www.vmware.com/partners/vmware/ESX_WAS_WP.pdf.
- "Scaling IBM DB2 9 in a VMware Infrastructure 3 Environment" at http://www.vmware.com/pdf/db2_scalability_wp_vi3.pdf.
- [Configuring VMware in a IBM Spectrum Control environment](#)
Before you can view information about VMware hypervisors and clusters, you must ensure that the VMware data sources are configured in IBM Spectrum Control.
- [VMware capacity reports](#)
You must probe both the ESX Server and the Storage Resource agent on the virtual machines before you can generate accurate reports for disk and file system

capacity.

Configuring VMware in a IBM Spectrum Control environment

Before you can view information about VMware hypervisors and clusters, you must ensure that the VMware data sources are configured in IBM Spectrum® Control.

About this task

To configure the VMware data sources in IBM Spectrum Control, complete the following steps:

Procedure

1. Add the VMware vSphere data source.
For example, add a vCenter Server, ESX, or ESXi hypervisor.
When you add a vCenter Server, each of the hypervisors that it manages are discovered and probed, as well as any clusters.
2. Schedule a probe to collect asset and status data about a hypervisor.
You must collect data before you can view and manage a hypervisor.
3. To manage the conditions that trigger alerts and the actions to take when those alerts are triggered, use the default alert policy that is automatically added to the hypervisor. Alternatively, you can configure alerts for each hypervisor so that you're notified of any critical conditions that are detected during a probe.

Related reference

- [Triggering conditions for hypervisor alerts](#)

VMware capacity reports

You must probe both the ESX Server and the Storage Resource agent on the virtual machines before you can generate accurate reports for disk and file system capacity.

About this task

The TOTAL row in the capacity report shows the capacity for the file system or the disk capacity. The total includes virtualized disks, virtual machines, non-virtualized disks, and non-virtualized machines.

For example, you have an ESX Server that has a capacity of 100 GB and 60 GB is allocated to the virtual machine. The virtual machine uses 5 GB of space. Both the ESX Server (H1) and the virtual machine (VC1) are probed. You also have a physical machine (PC1) that is probed. The capacity is shown in the report, as follows:

Column heading	Capacity	Used Capacity (calculated as Capacity minus Free Space)	Free space
TOTAL	130 GB	25 GB	105 GB
H1	100 GB	60 GB	40 GB
VC1	60 GB	5 GB	55 GB
PC1	30 GB	20 GB	10 GB

If you probed the physical machine (PC1) and the virtual machine (VC1) but did not probe the ESX Server (hypervisor), the capacity is shown as follows:

Column heading	Capacity	Used Capacity (calculated as Capacity minus Free Space; negative values are shown as 0)	Free Space
TOTAL	30 GB	0 GB	65 GB
PC1	30 GB	20 GB	10 GB
VC1	60 GB	5 GB	55 GB

If you probed the hypervisor (H1) and the physical machine (PC1) but did not probe the virtual machine (VC1), the capacity is shown as follows:

Column heading	Capacity	Used Capacity (calculated as Capacity minus Free Space)	Free Space
TOTAL	130 GB	80 GB	50 GB
H1	100 GB	60 GB	40 GB
PC1	30 GB	20 GB	10 GB

If you have an ESX Server that has devices that are directly attached to the virtual machine, the file system capacity created on the virtual machine is not added in the TOTAL row.

For example, the file system capacity shown in the table is based on the following factors:

- The file system (FS1) created on the hypervisor has a capacity of 100 GB and uses 60 GB of space.
- The file system (FS1) created from internal storage on the hypervisor (H1) has a capacity of 60 GB and uses 5 GB of space.
- The file system (FS2) on the device that is directly attached to the virtual machine (VC1) has a capacity of 5 GB and uses 4 GB of space.
- The file system created on the physical machine (PC1) has a capacity of 30 GB and uses 20 GB of space.

Column heading	Capacity	Used Capacity (calculated as Capacity minus Free Space)	Free Space
TOTAL	130 GB	24 GB	106 GB
H1 FS1	100 GB	60 GB	40 GB
VC1 FS1	60 GB	5 GB	55 GB
VC1 FS2	5 GB	4 GB	1 GB
PC1 FS1	30 GB	20 GB	10 GB

Planning for files systems and volume managers

Information about the file system formats and volume managers that are supported by IBM Spectrum® Control can help you set up and configure your environment.

About this task

- [File systems](#)
The file systems that are supported for monitoring and reporting by IBM Spectrum Control are listed.
- [Networked file systems](#)
This section describes support for IBM® General Parallel File System.
- [Volume managers](#)
This section lists the volume managers supported for monitoring by IBM Spectrum Control.

File systems

The file systems that are supported for monitoring and reporting by IBM Spectrum® Control are listed.

About this task

The Data Server supports monitoring and reporting of the following file systems:

- AIX® JFS, JFS2
 - Data ONTAP V7 for Network Appliance, including flexible volumes (FlexVol). With Flexvol, you can create multiple flexible volumes on a large pool of disks.
 - EXT2, EXT3, EXT4
 - FAT, FAT32
 - General Parallel File System (GPFS). For a list of the GPFS versions that are supported, view the information in the Agents, Servers and Browsers section of the [IBM Spectrum Control interoperability matrix](#).
 - HP_HFS
 - NTFS4, NTFS5
 - REISERFS
 - SAN File System (SANFS)
 - Temporary File System (TMPFS)
 - UFS
 - VMFS
 - Veritas File System (VxFS) r4.5 and 5 (all releases) on the following operating systems:
 - AIX
 - Linux®
- Tip: You can use IBM Spectrum Control with file systems and volumes that are created by using Veritas File System (VxFS) and Veritas Volume Manager (VxVM) releases 4 and 5, if the file systems and volumes conform to the features in Version 3 of VxFS and VxVM. IBM Spectrum Control cannot work with volumes and file systems with features that were introduced in VxFS and VxVM releases 4 and 5.
- WAFL
 - IBM Spectrum Control SAN File System

Networked file systems

This section describes support for IBM® General Parallel File System.

About this task

When using the Data server, you can use the monitoring and reporting of the IBM General Parallel File System (GPFS) 3.2 on AIX®.

- The IBM Spectrum® Control Storage Resource agent must be installed on one node within a GPFS *nodeset*. A nodeset is a collection of computers that see the same file system. If multiple agents are installed per GPFS nodeset, the first agent that sees the file system owns the file system. You cannot change the owning (or scanning) agent. If the agent is deleted, another agent takes ownership.

Because only one agent owns the file system and a file system cannot be scanned by more than one agent, there is no benefit to having more than one agent.

If more than one agent is installed, the last agent that runs a probe job takes ownership of the volume group. If a GPFS file system exists in that volume group and the agent that owns the volume group is not the agent that owns the file system, the file system information is not displayed in the volume group. The volume group is part of the asset tree. However, when the agent that owns the file system probes again, the data is corrected.

There must be physical access to the GPFS disk for IBM Spectrum Control to gather hardware disk information.

Volume managers

This section lists the volume managers supported for monitoring by IBM Spectrum® Control.

About this task

The Data Server supports the monitoring of the following volume managers:

- Veritas Volume Manager is supported on the following operating systems:
 - AIX®
 - Linux®
- AIX Logical Volume Manager (LVM)

Using these volume managers, you can create groups of logical volumes and disks. You can generate various reports for these disk and volume groups.

Planning for PowerHASystemMirror for AIX

IBM Spectrum® Control supports Storage Resource agents that are installed on IBM® PowerHA® SystemMirror® for AIX® nodes. Use this information to configure the PowerHA SystemMirror for AIX environment before you use it with IBM Spectrum Control.

Cluster resource groups

A *cluster resource group* represents a PowerHA SystemMirror for AIX entity.

You can configure PowerHA SystemMirror for AIX into cluster resource groups so that they can be highly available. You can define resource group policies and attributes that dictate how PowerHA SystemMirror for AIX manages resources to keep them highly available at different stages of cluster operation (startup, failover, and fallback). You can put the following types of resources into clustered resource groups.

Volume groups

A set of physical volumes that AIX treats as a contiguous, addressable disk region.

Logical volumes

A set of logical partitions that AIX makes available as a single storage unit. The logical volume is the "logical view" of a physical disk.

File system

A file system is written to a single logical volume. Typically, you organize a set of files as a file system for convenience and speed in managing data.

Shared file systems

A journaled file system that is entirely in a shared logical volume.

Applications

Critical services that are accessible to users.

Service IP labels or addresses

A way to establish communication between client nodes and the server node. Services, such as a database application, are provided by using the connection mode over the service IP label.

Tape resources

You can configure a SCSI or a Fibre Channel tape drive as a cluster resource in a non-concurrent resource group, making it highly available to two nodes in a cluster.

Communication links

You can define communication links as resources in a PowerHA SystemMirror for AIX resource group.

Planning for PowerHA SystemMirror for AIX support

With PowerHA SystemMirror for AIX, you can use the following types of configurations:

- PowerHA SystemMirror for AIX with cluster resources groups that are nonconcurrent. Concurrent cluster resource groups are not supported.
- The following volume groups in a PowerHA SystemMirror for AIX environment:
 - Standard volume groups
 - Enhanced concurrent-mode volume groups
 - Scalable volume groups

When you create an IBM Spectrum Control configuration, include the Storage Resource agent. A Storage Resource agent must be installed on each node of the cluster. All agents in a cluster must be configured to use the same listening port. You must also have a cluster resource group (CRG) with at least one IP address that is accessible from the IBM Spectrum Control server if you plan to run scans on the cluster resource group. Each Storage Resource agent collects information about the local node and the clustered resources that are currently hosted on the node.

The following operations are not supported.

- Probing and scanning of clustered databases
- Reporting on cluster information when the cluster is configured to have concurrent cluster resource groups

Note:

- Cluster resource group scanning is only supported by the IBM Spectrum Control Advanced Edition license.
- If the Storage Resource agent port is behind a firewall, that port must be opened for all virtual server and cluster resource group addresses.

Reports

The asset reporting navigation tree includes a By Cluster report that shows the nodes and cluster resource groups in each cluster. The disks, volume groups, and file systems that are shared as clustered resources are associated with the cluster resource group and not with the node where they are hosted.

When generating reports, you can use the By Cluster report subtype wherever the By Computer subtype is available. The By Cluster report summarizes the results of the nodes and the cluster resource groups in the cluster.

Monitoring groups

When you create a monitoring group or create a data collection schedule, you can select individual nodes and cluster resource groups. You can create a file system monitoring group that includes clustered and non-clustered resources.

Computer alerts

You can register for cluster-specific alerts that notify you that your cluster went through a configuration change. A cluster resource that is being added or removed is an example of such changes. You can also register for an alert that is triggered when a cluster resource group is moved. The same alert is used for a cluster resource group failover, cluster resource group fallback, and when the cluster resource group is manually moved to a new node. Cluster alerts will not be triggered until the next time the cluster is probed.

Installing PowerHA SystemMirror for AIX support

To install PowerHA SystemMirror for AIX support, complete the following steps:

1. Ensure that your PowerHA SystemMirror for AIX environment is configured for IBM Spectrum Control support.
2. Create a service IP label for the CRG with an IP address that can be contacted from the IBM Spectrum Control server.
3. Install a Storage Resource agent on each node of the cluster. All agents in the cluster must be configured to use the same listening port.
4. Run a discovery job for the agent.
5. Run a probe job on the agents.
6. Run a scan job on the agents.
7. Run a scan job on the cluster resource groups.
8. View reports for the PowerHA SystemMirror for AIX entities.
9. View the PowerHA SystemMirror for AIX environment in the topology viewer.

- [PowerHA SystemMirror for AIX environment](#)
Monitoring an IBM PowerHA SystemMirror for AIX environment with IBM Spectrum Control.

PowerHA SystemMirror for AIX environment

Monitoring an IBM® PowerHA® SystemMirror® for AIX® environment with IBM Spectrum® Control.

About this task

IBM Spectrum Control supports Storage Resource agents installed on PowerHA SystemMirror for AIX nodes. The Storage Resource agent must be installed on every node in the cluster. You cannot configure the Storage Resource agent as a clustered application.

The Storage Resource agent collects information from the cluster when the node is probed. Resources that are not clustered are reported under the node. The resources that are associated with a cluster resource group are reported under a computer entity that represents the cluster resource group. r.

The following entities are reported under the associated clustered resource group and not the node:

- Volume groups
- NFS shares
- Service IP labels

Physical volumes, logical volumes, and file systems for clustered volume groups are also reported under the cluster resource group.

Restriction: You can use PowerHA SystemMirror for AIX only with cluster resource groups that are nonconcurrent. You cannot use PowerHA SystemMirror for AIX with concurrent cluster resource groups.

The following information lists the support requirements for a PowerHA SystemMirror for AIX cluster environment.

Agent requirement

Each node in a PowerHA SystemMirror for AIX cluster must have a Storage Resource agent installed. IBM Spectrum Control can only monitor and report on nodes with a Storage Resource agent that is run in daemon mode. All Storage Resource agents that are installed in a cluster must listen on the same port number. Storage Resource agents in other clusters can be configured with a different port address. However, all agents in the cluster must use the port address used by the other agents in that cluster.

Probe requirements

Probes are not automatically executed in response to cluster events. Schedule probes to run as appropriate for the needs of the environment.

NAS support

You can use the information to plan your Network Attached Storage (NAS) support.

About this task

- [Network-attached storage system requirements](#)
Before adding network-attached storage (NAS) systems and NetApp devices for monitoring, ensure that you understand the system requirements.

Network-attached storage system requirements

Before adding network-attached storage (NAS) systems and NetApp devices for monitoring, ensure that you understand the system requirements.

NetApp devices

A NetApp device (7-mode only) can be added to IBM Spectrum® Control as a block storage device, a file storage device, or both.

- When you add a NetApp device as a block storage system, IBM Spectrum Control collects data by using a CIM agent.
- When you add a NetApp device as a file storage system, IBM Spectrum Control collects data through the native API (NAPI) of the device.

For a list of NetApp devices that can be monitored, see [NetApp support page](#).

Note: For new IBM Spectrum Control installations, Storage Resource agents are no longer supported for managing NetApp devices. If you upgrade from a previous version of IBM Spectrum Control, you can continue to manage NetApp devices that have Storage Resource agents installed from a previous release.

Other NAS devices

NAS devices other than Network Appliance devices must meet the following criteria in order to be supported by IBM Spectrum Control:

- If scanned from a UNIX agent, a NAS device must support Network File System (NFS) queries.
- If scanned from a Windows agent, a NAS device must support Common Internet File System (CIFS) queries.
- A NAS device must support SNMP queries (sysName and sysOID).
- A NAS device must supply a unique sysName.

The following conditions might be necessary for adding other NAS systems for monitoring and management by IBM Spectrum Control.

Agent requirements

You must install the agent on a machine that has access to the NAS filers within your environment that you want to monitor.

Windows

The agent that is logging in to and scanning the NAS filer is not required to be in the same domain as the user or the NAS filer. If you install the agent on a different domain from the NAS filer, the agent scans the NAS filer *if* the domain of the agent computer is a "trusted domain" by the domain of the NAS filer.

UNIX and Linux®

The agent computer must import the NAS filer export files as NFS mounts (or automounts on Oracle Solaris).

Note:

- You do not install agents to the NAS filers themselves.
- For IBM Spectrum Control 5196 Network Attached Storage 300 G machines only, install the agent directly on those machines. If the 300 G is clustered, you need to install an agent on each local node.

NAS requirements

The NAS filers within your environment must be visible to the machines where you install the agent or agents. If you want to monitor NAS filers from Windows, you must configure those NAS filers to be members of a Windows domain.

Microsoft Cluster Server

IBM Spectrum® Control can monitor and report on Microsoft Cluster Server (MSCS) clustered nodes and cluster resource groups.

About this task

Microsoft Cluster Server (MSCS) is a built-in feature of the Windows operating system. It is software that supports the connection of up to eight servers into a "cluster" for higher availability and easier manageability of data and applications. MSCS can automatically detect and recover from server or application failures. It can be used to move server workload to balance utilization and to provide for planned maintenance without downtime.

As with standard Windows clusters, you perform most of the configuration tasks, and also the management tasks, associated with Exchange clusters using the Cluster Administrator. Cluster Administrator is installed by default on servers that have Cluster Service installed.

You can also use Cluster Administrator to remotely administer a server cluster. Computers that are used to administer a server cluster remotely must be secure and restricted to trusted personnel.

When a cluster node is probed, there are normal errors in the logs for the disks belonging to the cluster resource groups that are not currently hosted on this node.

When looking at an MSCS cluster problem, the current state of the clustering environment is often needed. Microsoft provides a tool to dump the MSCS cluster configuration into a set of files.

Note: The term "virtual servers" is replaced with "cluster resource groups" in IBM Spectrum Control. Microsoft still refers to "virtual servers."

- [Microsoft Cluster Server environment](#)
To monitor resources in an MSCS environment, you must be aware of IBM Spectrum Control requirements.
- [Microsoft Cluster Server support](#)
Plan how to use IBM Spectrum Control to monitor resources in an MSCS environment.

Microsoft Cluster Server environment

To monitor resources in an MSCS environment, you must be aware of IBM Spectrum® Control requirements.

About this task

In an MSCS environment, you can collect information about local resources for a clustered node and the resources in a cluster resource group. The Storage Resource agent for IBM Spectrum Control is not cluster aware, so it cannot run in the cluster resource group and fail over from one node to another. However, the agent supports failover of a clustered file system so that IBM Spectrum Control probes can continue to work when the clustered file system is moved from one node to another.

The Storage Resource agent that runs on the node that hosts the cluster resource group is that agent discovers the resources on that cluster resource group.

Restriction: You cannot use IBM Spectrum Control to monitor a clustered database application.
To use IBM Spectrum Control in an MSCS cluster environment, consider the following requirements:

Server requirement

Not applicable. You cannot install the IBM Spectrum Control server on a node in an MSCS cluster.

Client requirement

Not applicable. You cannot install the IBM Spectrum Control client on a node in an MSCS cluster.

Agent requirements

- You must install a Storage Resource agent on each node in an MSCS cluster. IBM Spectrum Control can monitor and report only on nodes where a Storage Resource agent is installed and configured to run in daemon mode.
- If the Storage Resource agent port is behind a firewall, that port needs to be opened for all virtual server and cluster resource group addresses.
- All agents that are installed in a cluster must use the same port number to communicate with IBM Spectrum Control. Agents in other clusters can be configured with a different port address.
Tip: All agents in a cluster must use the port address used by the other agents in that cluster.

Microsoft Cluster Server support

Plan how to use IBM Spectrum® Control to monitor resources in an MSCS environment.

Cluster resource groups

A *cluster resource group* represents an MSCS entity.

You can configure MSCS into cluster resource groups. You can define resource group policies and attributes that dictate how MSCS manages resources to keep them highly available at different stages of cluster operation (startup, failover, and fallback). You can put the following types of resources into clustered resource groups.

Logical volumes

A set of logical partitions that MSCS makes available as a single storage unit. The logical volume is the "logical view" of a physical disk.

File system

A file system is written to a single logical volume. Typically, you organize a set of files as a file system for convenience and speed in managing data.

Shared file systems

A file system that is entirely in a shared logical volume.

Physical disk

A physical disk.

Network name

A network name.

IP address

An IP address.

Upgrading agents in an MSCS cluster

To upgrade the IBM Spectrum Control agents in an MSCS cluster, complete the following steps:

- Ensure that all Storage Resource agents in the cluster are configured to use the same port address.
- In the IBM Spectrum Control GUI, upgrade the agents.
- Do not probe agents in the cluster while the upgrade is in progress.
- When all agents in the cluster are upgraded successfully, run a probe on the cluster.

For more information about upgrading IBM Spectrum Control agents, see [Upgrading Storage Resource agents](#).

Configuring MSCS for monitoring

To use IBM Spectrum Control to monitor MSCS, complete the following steps:

1. Ensure that your MSCS environment is configured for IBM Spectrum Control support.
2. Install a Storage Resource agent on each node of the cluster. All agents in the cluster must be configured to use the same listening port.
3. Probe the nodes.
4. View information about the nodes on the Servers page.

Planning for the Virtual I/O Server

You can use the Storage Resource agent to gather information about Virtual I/O Servers. Before you can monitor Virtual I/O Servers, you must plan on how to install the agents in your environment.

Overview

The Virtual I/O Server is part of the IBM® PowerVM® hardware feature. The Virtual I/O Server allows the sharing of physical resources between LPARs including Virtual SCSI and virtual networking. Sharing of physical resources allows more efficient utilization of physical resources through sharing between LPARs and facilitates server consolidation.

Support for Virtual SCSI environment

The Virtual SCSI environment consists of a Virtual SCSI Server Adapter or Adapters that are created on the Virtual I/O Server and mapped to a Virtual I/O Client. If you have a Virtual SCSI environment and the virtual target device has a one to one mapping to a storage volume, IBM Spectrum® Control can extract the storage system serial number and correlate the Virtual SCSI disk on the LPAR to a corresponding storage volume. This function means that the Storage Resource agent installed on the Virtual I/O Client (LPAR), can collect this information. The Linux® system must be running on the IBM Power® System.

Note:

- IBM Spectrum Control performs the correlation only if the virtual target device is a one to one mapping to a single storage system volume.

The storage systems supported are:

- DS8000®
- SAN Volume Controller
- Storwize® V7000
- Storwize V7000 Unified

After the Storage Resource agent is installed on the Virtual I/O Server, you can see the following information:

- Storage system serial number
- Virtual SCSI Client Adapter
- Virtual SCSI disk
- Volume groups, logical volumes, and file system

Note:

- Currently, the correlation of Virtual SCSI disks that originate from the XIV® is not performed.
- IBM Spectrum Control can perform the correlation of Virtual SCSI disks only if AIX® Multiple Path I/O (MPIO) or SDDPCM multipath driver is used on the Virtual I/O Server.
- If there is one or more Virtual SCSI disks present on an LPAR, IBM Spectrum Control considers the LPAR as a virtual machine. In this case, IBM Spectrum Control does not include disk space for this LPAR in the TOTAL disk space reports. Disks that are assigned directly through the physical Fibre Channel adapter to the LPAR in a Virtual SCSI environment are listed as virtual disks in IBM Spectrum Control.

Login requirements

Use the **padmin** user ID when logging in to a Virtual I/O Server, **padmin** is the main administrator ID for Virtual I/O Servers. Upon login, a password change is required. There is no default password to remember.

The Virtual I/O Server does not support users logging in as root to install or configure IBM Spectrum Control agents. However, you can debug errors generated by agents as a root user. For example, you can run the **service.sh** script to gather debugging information.

Upon logging in to the Virtual I/O Server, you are placed into a restricted Korn shell. The restricted Korn shell works the same way as a regular Korn shell, except you cannot:

- Change the current working directory.
- Set the value of the SHELL, ENV, or PATH variable.
- Specify the path name of the command that contains a redirect output of a command with a >, >|, <>, or >.

As a result of these restrictions, you cannot run commands that are in locations not defined in your PATH variable. These restrictions prevent you from directly sending the output of the command to a file, instead requiring you to pipe the output to the **tee** command.

Installation requirements

The installation requirements are:

- IBM Spectrum Control supports Virtual I/O Server 1.5.2 or later.
- The Virtual I/O Server for POWER6® processor-based systems, must have a minimum of 768 MB of memory.

For information about disk space requirements for the Storage Resource agent, see [Hardware requirements](#) (Hardware requirements for the IBM Spectrum Control agents).

General procedure to install and configure agents

To install and configure the agents on the Virtual I/O Server, complete the following steps:

1. Install IBM Spectrum Control.
2. Install and configure the agents on the Virtual I/O Server.
3. When you install the agents, they are automatically started
4. Collect information about a Virtual I/O Server by running probe jobs.
5. View the storage information gathered by the monitoring jobs through the reports that you can generate through the IBM Spectrum Control GUI.

See the *IBM Spectrum Control User's Guide* for more information about how to use the IBM Spectrum Control user interface to collect and view information that is gathered by agents on the Virtual I/O Servers.

Planning to monitor Db2

It is a good practice to monitor your IBM® Db2® environment to better understand what is happening inside your Db2 data server. Db2 11.5 or later includes enhancements that make monitoring Db2 database environments more comprehensive with higher granularity of control.

About this task

This infrastructure is a superior alternative to the existing system monitor, event monitors, snapshot commands, and snapshot SQL interfaces.

For more information about monitoring your Db2 environment, see [IBM DB2® 11.5 for Linux®, Unix and Windows](#).

Sudo command privileges

When you install, operate, or uninstall IBM Spectrum® Control as a non-root user on AIX® or Linux® operating systems, you are required to have sudo privileges to multiple commands. Use command aliases in the /etc/sudoers file to configure your sudo privileges for multiple commands.

Example command aliases

The following command aliases are examples that you can create for installing, operating, and uninstalling IBM Spectrum Control:

- `Cmnd_Alias IBMSC_INSTALL = /usr/bin/xauth, /Downloads/Db2_11.5_image/server_dec/db2setup, /opt/ibm/db2/V11.5/cfg/db2ln, /Downloads/IBM_Spectrum_Control_image/SC/setup.bin`
- `Cmnd_Alias IBMSC_OPERATION = /opt/IBM/TPC/scripts/*.sh, /opt/IBM/TPC/scripts/ldap/*.sh, /opt/IBM/TPC/agent/bin/agent.sh, \ /opt/IBM/TPC/service/*.sh, /opt/IBM/TPC/service/changepasswords, /opt/IBM/TPC/client_images/*.sh, \ /opt/IBM/TPC/cli/*.sh, /opt/IBM/TPC/data/scripts/logCleanup.sh, /opt/IBM/TPC/data/scripts/EncryptPassword.sh, \ /opt/IBM/TPC/data/server/tools/repocopy, /opt/IBM/TPC/data/server/tools/runTPCDBMaintenance, \ /opt/IBM/TPC/device/bin/linux/setenv.sh, /opt/IBM/TPC/device/bin/linux/enableapialerts.sh, \ /opt/IBM/TPC/device/bin/linux/disableapialerts.sh, /opt/IBM/TPC/wlp/bin/server, \ /opt/IBM/TPC/wlp/bin/securityUtility, /opt/IBM/TPC/jre/bin/ikeyman, /opt/IBM/TPC/jre/bin/ikeycmd, /bin/kill`
- `Cmnd_Alias IBMSC_UNINSTALL = /opt/IBM/TPC/_uninst/uninstall`

The following examples show how to use the command aliases to provide non-root users privileges to the various commands:

```
username1 ALL=(root) SETENV: IBMSC_INSTALL
username2 ALL=(root) SETENV: IBMSC_OPERATION
username3 ALL=(root) SETENV: IBMSC_UNINSTALL
```

Installing

Before you install IBM Spectrum® Control, you must install IBM® Db2® because Db2 is required for the IBM Spectrum Control database repository.

- [IBM Spectrum Control components](#)
You can install IBM Spectrum Control in single-server or multiple-server environments. In a single-server environment, all components are installed on one server.
- [Hardware requirements](#)
The hardware requirements of the IBM Spectrum Control server.
- [Software requirements](#)
This section describes the software that is required to install and run your system. This section includes operating systems supported, browsers supported, databases supported, and other software required such as Db2.
- [Installing IBM Spectrum Control](#)
Install IBM Spectrum Control and its components in your environment. You can use the installation program or the command line in silent mode.
- [Installing IBM Cognos Analytics](#)
You can install the optional IBM Cognos® Analytics 11.2.0 or later reporting tool on one computer, on multiple servers for a distributed installation, or you can expand an existing single computer installation to another server to improve performance. Cognos Analytics provides an installation program that guides you through the installation process.
- [Configuring IBM Cognos Analytics](#)
After you install the optional IBM Cognos Analytics, to view predefined reports and create custom reports about IBM Spectrum Control, you must copy Db2 files, create a content store database, save your configuration settings, and start the services.
- [Reinstalling the software if a failure occurs](#)
If an installation failure occurs, you do not have to uninstall components that were successfully installed. IBM Spectrum Control provides an option to partially, or fully, roll back the installation.
- [Taking the first steps after installation](#)
After IBM Spectrum Control is installed, configure it to monitor the resources in your environment.

IBM Spectrum Control components

You can install IBM Spectrum® Control in single-server or multiple-server environments. In a single-server environment, all components are installed on one server.

In single-server or multiple-server environments, when you install IBM Spectrum Control, the following components are installed:

- Database repository
- IBM Spectrum Control servers, which comprise the following components:
 - Data server
 - Device server
 - Alert server
 - Export server
 - Web server
 - IBM Spectrum Control GUI

- Command-line interface (CLI)
- Storage Resource agent

IBM Spectrum Control in an IPv4 or IPv6 environment

IBM Spectrum Control can use both IPv4 and IPv6 addresses for communication between its components.

If you have a system that is configured for dual-stack networking (with both IPv4 and IPv6 addresses), IBM Spectrum Control defaults to IPv4 addressing.

For information about using IPv4 and IPv6, see [Planning for Internet Protocol Version 6](#).

Hardware requirements

The hardware requirements of the IBM Spectrum® Control server.

Important: IBM Spectrum Control server can require a large amount of memory, disk space, network bandwidth, and processor resources. To promote consistent availability and performance, install IBM Spectrum Control on a dedicated server where no other critical business applications, including IBM® Copy Services Manager, are installed. IBM Software Support is not responsible for availability, performance, and functional problems that are caused by other business applications that are deployed on the same server as IBM Spectrum Control.

For information about the hardware requirements for IBM Spectrum Control, see *Hardware Support: Memory, Processor and Disk Space* at [IBM Spectrum Control - Hardware Support: Memory, Processor and Disk Space](#).

For information about the components that can be used with IBM Spectrum Control, see *Find the Supported Hardware, Products and Platforms Interoperability Matrix Links* at <https://www.ibm.com/support/pages/node/388393>.

Restriction: On AIX® or Linux® operating systems, when you use a Virtual Network Computing (VNC) client that is connected to a VNC server, do not use the following key combinations:

- Alt + H to start the Help page
- Alt + N to proceed to next page
- Alt + P to return to the previous page
- Alt + I to start an installation action on the Summary page

You must press Tab and then press Enter or the spacebar when the button you want to press is selected. For example, to start the Help, press Tab until the Help button is highlighted and then press Enter or the spacebar.

Software requirements

This section describes the software that is required to install and run your system. This section includes operating systems supported, browsers supported, databases supported, and other software required such as Db2®.

About this task

- [Software requirements for operating systems](#)
The IBM Spectrum Control family supports a variety of operating systems.
- [Software requirements for Storage Resource agents](#)
Storage Resource agents can be used by IBM Spectrum Control in various operating systems.
- [Software requirements for the database repository](#)
IBM Spectrum Control uses Db2 as the database repository.
- [Web browser support](#)
The IBM Spectrum Control GUI runs in a web browser and includes the ability to monitor, manage, and troubleshoot storage resources. You can access this interface from anywhere that you have a web browser and connectivity to a network.
- [Software requirements for LDAP servers](#)
The LDAP repositories that are available on Windows, Linux® and AIX® operating systems, can be used with IBM Spectrum Control.
- [Software requirements for CIM agents](#)
IBM Spectrum Control requires CIM agents for some IBM storage devices.

Software requirements for operating systems

The IBM Spectrum® Control family supports a variety of operating systems.

Before you install IBM Spectrum Control, check the IBM Spectrum Control support site at <https://www.ibm.com/support/pages/node/6249361#Server> for the latest operating system support.

Software requirements for Storage Resource agents

Storage Resource agents can be used by IBM Spectrum® Control in various operating systems.

For a list of the operating systems that are supported for Storage Resource agents, go to <https://www.ibm.com/support/pages/node/6249361#Agents>.

Software requirements for the database repository

IBM Spectrum® Control uses Db2® as the database repository.

Important: You must use a 64-bit version of Db2 with IBM Spectrum Control.

Supported versions of Db2

For a list of the supported Db2 versions, go to <http://www.ibm.com/support/pages/node/6249361#DB>.

Web browser support

The IBM Spectrum® Control GUI runs in a web browser and includes the ability to monitor, manage, and troubleshoot storage resources. You can access this interface from anywhere that you have a web browser and connectivity to a network.

Before you start IBM Spectrum Control, ensure that you are using a supported web browser. For a list of web browsers that you can use with IBM Spectrum Control, see the support matrix at [IBM Spectrum Control - Platform Support: Servers, Agents, and Browsers - Web Browsers](#).

Note: The IBM Spectrum Control GUI requires the use of cookies to manage user preferences, such as table customization and frequently viewed charts. Ensure that cookies are enabled to use these features.

Software requirements for LDAP servers

The LDAP repositories that are available on Windows, Linux® and AIX® operating systems, can be used with IBM Spectrum® Control.

For information about the LDAP repositories that are supported, see one of the following WebSphere® Application Server Liberty web pages:

- For version information, see [WebSphere Liberty - Detailed Requirements](#) - On the Supported Software tab, click LDAP Servers.
- For fix pack information, see [Fix list for IBM® WebSphere Application Server Liberty](#).

Software requirements for CIM agents

IBM Spectrum® Control requires CIM agents for some IBM® storage devices.

CIM agents for switches and directors

CIM agents (SMI-S providers) are also required for switches and directors. For a complete list of switches and directors that are supported and the level of CIM agents that are required for each device, see <https://www.ibm.com/support/pages/node/388393>.

CIM agents for non-IBM storage systems

IBM Spectrum Control also supports some non-IBM storage systems. For example, IBM Spectrum Control supports HP, TagmaStore, CLARiiON, Symmetrix, and other storage systems.

For a complete list of storage systems that are supported and the level of CIM agents that are required for each device, see <https://www.ibm.com/support/pages/node/388393>.

Installing IBM Spectrum Control

Install IBM Spectrum® Control and its components in your environment. You can use the installation program or the command line in silent mode.

- **[Installation checklists for IBM Spectrum Control](#)**
The installation of IBM Spectrum Control requires you to install a number of components. You can use the following checklists to guide you through the installation procedures for each component. Use separate checklists for multiple server installations.
- **[Db2](#)**
You can install the most recent version of Db2 that IBM Spectrum Control supports on the Linux®, AIX®, or Windows operating systems.
- **[Starting the installation programs](#)**
There are various methods to start the IBM Spectrum Control installation programs and on various operating systems.
- **[Installing IBM Spectrum Control in a single-server environment](#)**
Install IBM Spectrum Control by using the installation program or the command line in silent mode.
- **[Installing IBM Spectrum Control in a multiple-server environment](#)**
You can install IBM Spectrum Control by using the installation program or by using silent mode from the command line. Note, installing IBM Spectrum Control by using console mode is not supported.
- **[Accessing IBM Spectrum Control server from a remote computer using Command Line Interface](#)**
Accessing client components on a remote computer helps you complete storage management tasks from a computer that does not have a IBM Spectrum Control server that is installed on it.
- **[Installing IBM Spectrum Control on a Windows domain](#)**
You can install IBM Spectrum Control and Db2 by using a Windows domain or a local user account.

- [Verifying the connection to the domain controller computer by using the Dcdiag tool](#)
Use the Dcdiag command line tool to help you determine whether the domain controller computer is registered with the domain name server (DNS), whether the controller can be pinged, and whether the controller has Lightweight Directory Access Protocol (LDAP) connectivity.
- [Installing IBM Spectrum Control and associated products using minimal space on the Windows C: drive](#)
In some cases, you might not want to install Db2 or IBM Spectrum Control on the C: drive.
- [Verifying the installation](#)
After you install IBM Spectrum Control, you can verify whether the installation was successful.
- [Reviewing the log files to resolve installation issues](#)
If an error occurs during IBM Spectrum Control installation, you can review the error log files to resolve issues and continue the installation.
- [Changing languages](#)
You can change the operating system language and web browser that determines the language in which IBM Spectrum Control services and the IBM Spectrum Control GUI are displayed.
- [Adding an installation license](#)
If you installed IBM Spectrum Control and want to add a IBM Spectrum Control license so that you can monitor storage systems with different license models, you can use the installation program or silent-mode installation.
- [Installing Storage Resource agents](#)
You can install Storage Resource agents by using the IBM Spectrum Control user interface or a command.

Installation checklists for IBM Spectrum Control

The installation of IBM Spectrum® Control requires you to install a number of components. You can use the following checklists to guide you through the installation procedures for each component. Use separate checklists for multiple server installations.

Prepare your environment for installation

Ensure that your target environment is ready for an installation.

Table 1. Prepare your environment for an installation check list

Check	Required for	Task
<input type="checkbox"/>	Antivirus software	<p>Installing IBM Spectrum Control involves making use of your operating system in manners typical for installing new application software. If your antivirus software is set on the maximum mode, it might prevent some of your changes from being accepted.</p> <p>To verify that your installation completes correctly, enable your antivirus software product to allow the following instances:</p> <ul style="list-style-type: none"> • The /etc/hosts file can be edited • Files can be created in the /temp directory • New executable files can be created in the C:\Program Files directory <p>If McAfee Adaptive Threat Protection is enabled on the server where IBM Spectrum Control is installed, it might prevent some services from starting or stopping. To help avoid this issue, open McAfee Adaptive Threat Protection and go to settings. In the Real Protect Scanning (Windows only) section, verify if Enable client-based scanning is selected. If so, select Low from the Sensitivity level list.</p> <p>For more information about McAfee Adaptive Threat Protection, see the following links:</p> <ul style="list-style-type: none"> • Adaptive Threat Protection — Options • Overview of Adaptive Threat Protection
<input type="checkbox"/>	Fully qualified host names	IBM Spectrum Control requires fully qualified host names, so you must install the product on a computer that has a fully qualified host name. For example, abc.tucson.example.com is a fully qualified host name where abc is the host name and tucson.example.com is the full domain name.
<input type="checkbox"/>	Hardware prerequisites	For information about the hardware requirements for IBM Spectrum Control, see <i>Hardware Support: Memory, Processor and Disk Space</i> at https://www.ibm.com/support/pages/node/388393 .
<input type="checkbox"/>	Software prerequisites	<ul style="list-style-type: none"> • The IBM Spectrum Control family supports various operating systems, see Software requirements for operating systems • For the database repository, see Software requirements for the database repository • For LDAP servers, see Software requirements for LDAP servers
<input type="checkbox"/>	Required user privileges	<p>Ensure that you have the required privileges to install Db2 and IBM Spectrum Control.</p> <p>If you are installing on a Windows operating system, you must install as a user with administrator privileges. If you are installing on an AIX or Linux operating system, you must install as a root user or as a non-root user who has sudo privileges.</p> <p>If you want to install as a non-root user, your system administrator must complete configuration tasks on the target server before you begin the installation. These tasks are listed in Installing Db2 on AIX or Linux and Installing IBM Spectrum Control in a single-server AIX or Linux environment.</p> <p>Ensure that Windows administrators have the Debug programs privilege.</p>
<input type="checkbox"/>	Download packages	<p>You must download the installation packages or images into a separate directory for each product. The following list provides suggested download directory paths:</p> <ul style="list-style-type: none"> • Windows operating systems: <ul style="list-style-type: none"> ◦ c:\downloads\DB2 ◦ c:\downloads\TPC • AIX or Linux operating systems: <ul style="list-style-type: none"> ◦ /downloads/DB2 ◦ /downloads/TPC <p>For more information, see Planning for installation</p>

Check	Required for	Task
<input type="checkbox"/>	Language setting	Determine the operating system language that you want to use for IBM Spectrum Control. For optimal support, it's recommended that you use English. The web browser determines the language in which IBM Spectrum Control services and the IBM Spectrum Control GUI are displayed. For more information, see Changing languages .

Windows domain

Before you can install IBM Spectrum Control on a Windows domain, you must determine which installation method is appropriate, based on your environment. For more information, see [Planning to install IBM Spectrum Control in a Windows domain](#)

Table 2. Windows domain installation check list

Check	Required for	Task
<input type="checkbox"/>	Windows domain and local user accounts	When a computer is a member of a Windows domain, you can install IBM® Db2® on the computer such that the Db2® installation process creates a local Db2® user account or a domain Db2 user account. n Important: Before you run your Db2 installation and begin your IBM Spectrum Control installation, please validate that the Windows server hostname is 15 characters or less in length. See Windows domain and local user accounts .
<input type="checkbox"/>	Adding a computer to the Windows domain	Before you can install IBM Spectrum Control in a Windows domain, you must first add the computer on which you plan to install IBM Spectrum Control to the domain, see Adding a computer to the Windows domain .
<input type="checkbox"/>	Create Windows domain common user account	You must create a Windows domain common user account before you can install IBM Spectrum Control in a Windows domain, see Creating a Windows domain common user account for IBM Spectrum Control .
<input type="checkbox"/>	Grant Db2 SYSADM authority to Windows domain user account	If a Windows domain user account is used to install IBM Spectrum Control, the user account might not have the Db2 SYSADM authority, because Db2 goes to the domain controller computer to list the groups. Before you install IBM Spectrum Control, you must grant the Windows domain user accounts the Db2 SYSADM authority, see Granting Db2 SYSADM authority to a Windows domain user account .

Db2

Install Db2 11.5.7 on Linux, AIX, or Windows operating systems.

Table 3. Db2 installation check list

Check	Required for	Task
<input type="checkbox"/>	Prepare for Db2 installation	<ul style="list-style-type: none"> When you install Db2 11.5.7, a default user name and groups are created. This default user name and password are used to install IBM Spectrum Control. See Preparing to install Db2.
<input type="checkbox"/>	Install on Windows	For information on how to install Db2 on Windows operating systems, see Installing Db2 on Windows .
<input type="checkbox"/>	Install Db2 by using a Windows domain user account	Before you install IBM Spectrum Control in a Windows domain, you must install Db2 and register the Db2 license key. You must have domain administrator access. Important: Before you run your Db2 installation please validate that the Windows server hostname is 15 characters or less in length. You do not need to complete this task if you are planning to install Db2 by using a local user account, see Installing Db2 by using a Windows domain user account .
<input type="checkbox"/>	Install Db2 on AIX or Linux - GUI installation	For information on how to install Db2 on AIX or Linux by using the GUI installation process, see Installing Db2 on AIX or Linux .
<input type="checkbox"/>	Install Db2 on AIX by using the command line	For more information on how to install Db2 on AIX by using the command line, see Installing Db2 on AIX or Linux by using the command-line .
<input type="checkbox"/>	Verify Db2 installation	You can verify that Db2 11.5.7 is installed properly by using the command line processor (CLP) or the First Steps GUI, see Verifying that Db2 is installed correctly .

IBM Spectrum Control

Install IBM Spectrum Control by using the installation program or the command line in silent mode.

Table 4. Installation check list

Check	Required for	Task
<input type="checkbox"/>	Choose installation method	For more information about what methods are available to install IBM Spectrum Control, see Planning for installation
<input type="checkbox"/>	Install IBM Spectrum Control on a single server.	For more information about how to install IBM Spectrum Control on a single server, see Installing IBM Spectrum Control in a single-server environment
<input type="checkbox"/>	Install IBM Spectrum Control on multiple servers.	For more information about how to install IBM Spectrum Control on multiple servers, see Installing IBM Spectrum Control in a multiple-server environment

Db2

You can install the most recent version of Db2® that IBM Spectrum® Control supports on the Linux®, AIX®, or Windows operating systems.

About this task

For more information about the supported versions, see [Software requirements for the database repository](#).

- **Preparing to install Db2**
Before you install Db2, run the **db2prereqcheck** command on the target installation server to ensure that the server meets the prerequisites for Db2. For example, **db2prereqcheck -v 11.5.7.0**. You can ignore the section on the Db2 pureScale® prerequisites.
- **Installing Db2**
You can install IBM® Db2 on AIX, Linux, or Windows operating systems.
- **Verifying that Db2 is installed correctly**
You can verify that Db2 is installed properly by using the command-line processor (CLP) or the First Steps GUI.
- **Licensing Db2**
You can apply a valid license to Db2 by using the **db2licm** command.

Preparing to install Db2

Before you install Db2, run the **db2prereqcheck** command on the target installation server to ensure that the server meets the prerequisites for Db2. For example, **db2prereqcheck -v 11.5.7.0**. You can ignore the section on the Db2 pureScale® prerequisites.

About this task

Before you install Db2, a default username and groups are created. This default username and password are used to install IBM Spectrum® Control. You can use the default username and password or provide your own.

Note: Windows operating systems do not support systems that are configured only for IP 6. (Windows is enabled for IP 6 and IP 4.)

- **Preparing to install Db2 on Windows**
Before you install Db2 on Windows operating systems, ensure that you correctly configured your directory names.
- **Preparing to install Db2 on UNIX or Linux**
To operate Db2 on operating systems such as UNIX or Linux, the instance owner, the fenced user, and the Db2 administration server user are required. These users and groups are automatically created when Db2 is installed.

Preparing to install Db2 on Windows

Before you install Db2 on Windows operating systems, ensure that you correctly configured your directory names.

Short file names

If Db2 is installed on a drive on which 8.3 file names are disabled, and the Db2 installation directory name has spaces in it, Db2 must be reinstalled on a drive that has 8.3 file names enabled, or in a directory that does not have spaces in the name.

To check the current 8.3 file name settings for a drive, in a Db2 command window, run the **'fsutil.exe behavior query disable8dot3'** command.

These are the values when you run the command:

- **0** = Create 8.3 short file names (default)
- **1** = Do not create 8.3 file names
- **2** = Set 8.3 file names on a per volume basis
- **3** = Disable 8.3 file names on all volumes, except the system volume

The registry state of **NtfsDisable8dot3NameCreation** is **0**, which means that 8.3 file names are enabled on all volumes.

To enable 8.3 file names on a global basis, run the **fsutil.exe behavior set disable8dot3 0** command.

Preparing to install Db2 on UNIX or Linux

To operate Db2® on operating systems such as UNIX or Linux®, the instance owner, the fenced user, and the Db2 administration server user are required. These users and groups are automatically created when Db2 is installed.

About this task

The user ID you need to install IBM Spectrum® Control must be the owner of the instance that you want to use; this user ID and password are created when you install Db2.

[Table 1](#) provides a list of the default user and group names.

Table 1. Db2 setup default user and group names

Required user	Default user name	Default group name	Description
instance owner	db2inst1	db2iadm1	Created in the instance owner home directory. This user ID controls all Db2 processes and owns all file systems and devices used by the databases contained within the instance.
fenced user	db2fenc1	db2fadm1	Used to run user-defined functions and stored procedures that are separate from the address space that is used by the Db2 database.

Related information

- 🔗 [Db2 User, user ID and group naming rules](#)
- 🔗 [Db2 password rules](#)

Installing Db2

You can install IBM® Db2® on AIX®, Linux®, or Windows operating systems.

- [Installing Db2 on Windows](#)
Before you can install IBM Spectrum Control, you must install Db2®.
- [Installing Db2 on AIX or Linux](#)
You can install Db2 as a root user or non-root user on an AIX or Linux operating system by using the Db2 Setup program. To install Db2 as a non-root user, the user must have sudo privileges on the target AIX or Linux server.
- [Installing Db2 on AIX or Linux by using the command-line](#)
You can install Db2 on AIX or Linux operating systems by using commands.

Installing Db2 on Windows

Before you can install IBM Spectrum® Control, you must install Db2®.

Before you begin

If you are installing Db2 in a Windows domain environment: See [Installing Db2 by using a Windows domain user account](#) for additional and alternate steps that are required for the installation of Db2.

Tip: You can disable the IBM® SSH server in Db2 11.1 or later because it is a Db2 service that is not used and can appear as an issue in security scans. See [IBM Spectrum Control Server Vulnerability Scan Reports Weak Cipher](#).

Note, Db2 11.5.5 or later on Windows does not include the IBM SSH Server.

About this task

To install Db2 on a Windows operating system, complete the following steps:

1. Log on with a user ID that has Administrator authority on Windows.
2. Use the extracted Db2 installation image. Windows Autorun starts the launchpad.
If you are using Windows Autorun, the installation program should start in 15-30 seconds. If the installation program does not start, complete one of the following steps:
 - In a command prompt, to start the Db2 set up page, go to `web_image_extraction\server_folder` and run the `setup.exe` command.
 - In Windows Explorer, go to `web_image_extraction\server_folder`, and double-click the setup.exe file.Where `server_folder` in the preceding steps is the location of the Db2 installation program.
3. In the DB2® Setup Launchpad, click Install a Product.
4. Click Install New for Db2 11.5.7.0 Server Editions.
Note: If you do not see the Db2 Setup program, it is probably running in the background. Look for a blinking icon in the Windows taskbar and click it to bring it to the foreground.
5. On the Welcome page of the Db2 Setup program, click Next.
6. On the Software License Agreement page, review and accept the license agreement, and click Next.
7. On the Select the installation type page, click Typical, and click Next.
8. On the Select the installation, response file creation, or both page, select Install Db2 Server Edition on this computer and save my settings in a response file.
9. Enter a response file name or accept the default, and click Next.
10. On the Select the installation folder page, enter a directory or accept the default, and click Next.
11. On the Set user information for the Db2 Administration Server page, enter the following user information:
 - a. Leave the Domain field blank.
 - b. In the User name field, type the Db2 user ID that you want to use or accept the default.
 - c. In the Password field, type a password and confirm the password by typing it again in the Confirm password field.
12. Select the Use the same account for the remaining Db2 services check box and click Next.
13. On the Configure Db2 instances page, click Next.
14. On the Set up notifications page, clear the Set up your Db2 server to send notifications check box, and click Next.
15. On the Enable operating system security for Db2 objects page, accept the defaults, and click Next.
The default is to enable operating system security.

Note: If you installed Db2 before on this system, and the DB2ADMS group exists, when you click Next, the following message is displayed:

Warning

The group name "DB2ADMS" already exists in the system and will be granted complete access to the DB2 folders and objects through the operating system.
Click OK to continue installation, or click Cancel to input another group name.

16. On the Start copying files and create response file page, review the current settings, and click Finish.
17. On the Setup is complete page, review the information, and click Next.
18. On the Install additional products page, click Finish.
19. The setup program closes and the Db2 First Steps page opens. Close the Db2 First Steps page.
20. Restart the system.

Note: To prevent the Db2 logs from filling up the C drive on the system you are installing on, you can set the *DIAGPATH* and *SPM_LOG_PATH* variables to point to an external, fast disk drive.

What to do next

After you restart your computer, you must install IBM Spectrum Control. For more information about installing IBM Spectrum Control in a single-server environment, see [Installing IBM Spectrum Control in a single-server Windows environment](#). For information about installing IBM Spectrum Control in a multiple-server environment, see [Installing IBM Spectrum Control with a remote database by using the installation program](#).

Installing Db2 on AIX or Linux

You can install Db2® as a root user or non-root user on an AIX® or Linux® operating system by using the Db2 Setup program. To install Db2 as a non-root user, the user must have sudo privileges on the target AIX or Linux server.

Before you begin

The X Window System is required on the AIX or Linux server to install Db2 by using the installation program.

If you want to install Db2 as a non-root user, have your system administrator complete the following tasks on the target server before you begin the installation:

- If you are installing Db2 on an AIX server, install sudo on the server. The default AIX distribution does not include sudo. To get the sudo installation package, go to <https://ibm.biz/BdftKL>. Find *sudo-version* in the Package column, and click RPM to download the package.
- Ensure that the non-root user has sudo privileges for the following commands on the server:
 - **xauth** (Enables the non-root user to provide access to their X Window System display for the root user.)
 - **db2setup** (Enables the non-root user to launch the Db2 installation program.)

You must enter the full path to the commands when you configure the sudo privileges. For example:

```
username      ALL=(root)      SETENV: /usr/bin/xauth,/home/username/Downloads/DB2/server_dec/db2setup
```

- Ensure that the following requirements are met to enable the installation of Db2 by using the Setup program:
 - The non-root user is able to use the X Window System on the target server and the X Window System **DISPLAY** environment variable is set correctly for the not-root user.
 - The **DISPLAY** environment variable is preserved in the sudo environment.
 - The root user has access to the X Window System display that is owned by the non-root user.
- If you are installing Db2 on an AIX server, ensure that the **ODMDIR** environment variable is preserved in the sudo environment.

Procedure

1. Log on to the target server as the root user or as a non-root user who has sudo privileges.
2. Use an extraction tool to extract the installation image files on the server.
3. In a command shell, change the directory to the location of the Db2 installation program db2setup, and then enter one of the following commands:

Root user

```
./db2setup
```

Non-root user

```
sudo ./db2setup
```

4. On the Db2 Setup Welcome page, click New Install.
5. On the Choose a Product page, select Db2 11.5.7.0 Server Editions and click Next.
6. On the Configuration page:
 - a. Enter an installation directory or accept the default.
 - b. Select Typical as the installation type.
 - c. Select Create an instance.
 - d. Review and agree to the IBM® terms.
 - e. Click Next.
7. On the Instance Owner page, select New user, and enter the following information:
 - a. User name
Db2 adds this user ID to the **db2iadm1** group, and if this user ID does not exist, Db2 creates it.
 - b. Group name
 - c. Password
 - d. Home directory
The Db2 instance owner user is the user that you enter when you install IBM Spectrum Control.
8. Select Use default UID and Use default GID, and click Next.
9. On the Fenced User page, select New user, and enter the following information:
 - a. User name
Db2 adds this user ID to the **db2fadm1** group, and if this user ID does not exist, Db2 creates it.
 - b. Group name
 - c. Password
 - d. Home directory
The fenced user is used to run user-defined functions and stored procedures that are separate from the address space that is used by the Db2 database.
10. Select Use default UID and Use default GID, and click Next.
11. On the Response File and Summary page:

- a. Select Install Db2 Server Edition on this computer and save my settings in a response file.
 - b. Enter a response file name or accept the default.
 - c. Review the Summary and click Finish.
12. On the Setup has completed successfully page, review the information, and click Finish.

What to do next

To ensure that Db2 starts when the server starts, run the **db2iauto** command.

Root user

```
DB2_installation_directory/bin/db2iauto -on DB2_instance_name
```

Non-root user

```
sudo DB2_installation_directory/bin/db2iauto -on DB2_instance_name
```

Where *DB2_installation_directory* is the location of Db2 and *DB2_instance_name* is the instance that you created when you installed Db2. The default instance name is *db2inst1*.

Installing Db2 on AIX or Linux by using the command-line

You can install Db2® on AIX® or Linux® operating systems by using commands.

Before you begin

Ensure that you have:

- 2 - 3 GB of hard disk drive space for the Db2 installation **tar** file and extracted files.
- A file system with at least 30 GB for the IBM Spectrum® Control repository.

Procedure

To install Db2 on AIX or Linux operating systems by using a command line, complete the following steps:

1. Log in with a user ID that has root authority.
2. Install Db2.
 - a. Create a temporary directory (for example, **db2temp**) to hold the Db2 installation image and extracted files.
 - b. Copy or download the Db2 installation image into **db2temp**. The name of the file varies depending on the location from where the file is downloaded or copied and the language to which it is associated.
 - c. Use an extraction tool to extract the installation image files.
For example, if the name of the file is **v11.5_aix64_server.tar.gz**, enter **tar -zxf v11.5_aix64_server.tar.gz**. In this example, **v11.5_aix64_server.tar.gz** is extracted in the **db2temp** directory.
 - d. Change directory to the location of the Db2 installation program: **db2_install**.
 - e. Enter **./db2_install** to run the command-line installation program.
 - f. Enter yes to accept the license agreement terms.
 - g. Enter yes to install into the default directory; enter no to install into a directory you specify.
 - h. Enter **SERVER** for the Db2 product to install.
 - i. Enter no because you do not want to install the Db2 pureScale® feature.
The Installation Summary is displayed, which indicates a completed installation. Errors related to the install of the TSAMP file set can be ignored.
3. Create users and groups for Db2 on AIX.
 - a. Enter the following commands to create the Db2 groups:

```
mkgroup db2iadml
```

```
mkgroup db2fadml
```

- b. Enter the following commands to create the Db2 users and assign them to the Db2 groups:

```
mkuser pgrp=db2iadml groups=db2iadml home=/home/db2inst1 db2inst1
```

```
mkuser pgrp=db2fadml groups=db2fadml home=/home/db2fenc1 db2fenc1
```

- c. Enter the following commands to set the correct ownership on the Db2 users' home directories:

```
chown -R db2inst1:db2iadml /home/db2inst1
```

```
chown -R db2fenc1:db2fadml /home/db2fenc1
```

- d. Enter the following commands to set the password for the Db2 user db2inst1:

```
passwd db2inst1
```

```
pwdadm -f NOCHECK db2inst1
```

- e. Enter the following commands to set the password for the Db2 user db2fenc1:

```
passwd db2fenc1
```

```
pwdadm -f NOCHECK db2fenc1
```

- f. Enter the following command to add the Db2 user db2inst1 to the groups staff and system:

```
usermod -G staff,system db2inst1
```

4. Create users and groups for Db2 on Linux:

a. Enter the following commands to create the Db2 groups:

```
groupadd db2iadm1
```

```
groupadd db2fadm1
```

b. Enter the following commands to create the Db2 users:

```
useradd db2inst1
```

```
useradd db2fenc1
```

c. Enter the following commands to assign the Db2 users to the Db2 groups:

```
usermod -a -G db2iadm1 db2inst1
```

```
usermod -a -G db2fadm1 db2fenc1
```

d. Enter the following commands to set the correct ownership on the Db2 users' home directories:

```
chown -R db2inst1:db2iadm1 /home/db2inst1
```

```
chown -R db2fenc1:db2fadm1 /home/db2fenc1
```

e. Enter the following command to set the password for the Db2 user db2inst1:

```
passwd db2inst1
```

f. Enter the following command to set the password for the Db2 user db2fenc1:

```
passwd db2fenc1
```

g. Enter the following command to add the Db2 user db2inst1 to the group root:

```
usermod -a -G root db2inst1
```

5. Create a Db2 instance named **db2inst1**:

a. Change directory to `<Db2_installation_dir>/instance`.

b. Enter the command:

```
./db2icrt -a server -u db2fenc1 db2inst1
```

c. Enter the command:

```
. /home/db2inst1/sqllib/db2profile
```

d. Enter the command:

```
db2start
```

6. Configure Db2 communication.

a. Examine the `/etc/services` file. If the line **db2c_db2inst1**

25010/tcp exists in that file, proceed to Step 7. If the line **db2c_db2inst1**
25010/tcp does **not** exist in that file, do the following:

b. Edit the `/etc/services` file and add the following line at the end of the file: **db2c_db2inst1 25000/tcp**.

c. Enter the command:

```
db2 update dbm cfg using svcename db2c_db2inst1
```

d. Enter the command:

```
db2set DB2COMM=tpcipc
```

e. Enter the command:

```
db2stop
```

f. Enter the command:

```
db2start
```

7. Enter the command: **db2level** to verify the installed version of Db2.

Verifying that Db2 is installed correctly

You can verify that Db2® is installed properly by using the command-line processor (CLP) or the First Steps GUI.

About this task

To verify that Db2 is installed, complete the following steps:

1. Create the SAMPLE database.
2. Connect to the SAMPLE database.
3. Run a query against the SAMPLE database.
4. Drop the SAMPLE database.

- [Verifying Db2 installation using the command-line processor \(CLP\)](#)
You can verify that Db2 is successfully installed using the command-line processor (CLP).
- [Verifying Db2 installation by using the First Steps tool](#)
You can verify that Db2 is installed successfully, by using the First Steps tool.

Verifying Db2 installation using the command-line processor (CLP)

You can verify that Db2® is successfully installed using the command-line processor (CLP).

About this task

To verify that Db2 is installed using the command line processor, complete the following steps:

Procedure

1. Log on to the system as a user with SYSADM authority.
2. Enter `db2start` to start the database manager.
In operating systems such as AIX® and Linux®, you must source the `db2profile` before you run the **db2start** command. For more information about sourcing the profile, go to the IBM Spectrum® Control Knowledge documentation and search for *Using the command line on AIX and Linux*.
3. Enter the **db2sample** command to create the **SAMPLE** database.
This command might take a few minutes to process. There is no completion message. When the command prompt returns, the process is complete. The **SAMPLE** database is automatically cataloged with the database alias **SAMPLE** when it is created.
4. Enter the following commands in a Db2 command window to connect to the SAMPLE database, retrieve a list of all the employees who work in Department 20, and reset the database connection:

```
db2 connect to sample
db2 "select * from staff where dept = 20"
db2 connect reset
```
5. After verifying the installation, remove the **SAMPLE** database to free up disk space.
For more information about verifying Db2 installation, refer to the Db2 documentation for your operating system.
6. Enter `db2 drop database sample` to drop the **SAMPLE** database.

Verifying Db2 installation by using the First Steps tool

You can verify that Db2® is installed successfully, by using the First Steps tool.

Before you begin

- Ensure that the domain user account that you use to create the sample database has SYSADM or SYSCTRL authority.
- Install IBM® Data Studio and the First Steps component.

About this task

The First Steps tool is part of the getting started component grouping in the Db2 setup program. It is installed as part of a typical installation or you can select the First Steps tool during a custom installation.

Procedure

To verify the Db2 installation by using the First Steps tool, complete the following steps:

1. Log on to the system with the user account that you want to use to verify the installation.
2. Start First Steps by completing one of the following steps:
 - On operating systems such as AIX® or Linux®, run the `db2fs` command.
 - On Windows, run the `db2fs.bat` command.
3. On the First Steps window, select Create Sample Databases.
4. On the Create Sample Databases page, select the databases that you want to create and click OK.
By default, the SAMPLE database is created on the computer where Db2 is installed.
5. Click OK.
6. Open Data Studio.
7. Expand the navigation tree to view the SAMPLE database and SAMPLE database objects.
8. Select the Tables object to view the SAMPLE database tables.
9. After you verify the installation, remove the SAMPLE database to free disk space.
10. Enter the **db2 drop database sample** command to drop the SAMPLE database.

Licensing Db2

You can apply a valid license to Db2® by using the **db2licm** command.

About this task

The installation and operation of IBM Spectrum® Control requires that Db2 is running on your system. Important: After your Db2 11.1 trial license period expires, Db2 11.1 does not start without a valid license. IBM Spectrum Control does not support the Db2 11.5 Community license.

Procedure

To apply a valid license to Db2 with the **db2licm** command, complete the following steps:

1. Retrieve and extract a valid Db2 license that is supplied on the IBM Spectrum Control download site.
2. Open a Db2 command prompt or switch to the Db2 instance owner user and in the location where you extracted your Db2 license, change to the following directory:

```
awse_o/db2/license/          if you are using Db2 11.1
std_vpc/db2/license/         if you are using Db2 11.5
```

3. Run the following command:

Using Db2 11.1 on the Windows operating system:

```
Db2_installation_location\BIN\db2licm -a db2awse_o.lic
```

Using Db2 11.5 on the Windows operating system:

```
Db2_installation_location\BIN\db2licm -a db2std_o.lic
```

where *Db2_installation_location* is the base installation directory for Db2 on your computer. For example, C:\Program Files\IBM\SQLLIB.

Using Db2 11.1 on the AIX®/Linux® operating system:

```
Db2_installation_location/adm/db2licm -a db2awse_o.lic
```

Using Db2 11.5 on the AIX/Linux operating system:

```
Db2_installation_location/adm/db2licm -a db2std_o.lic
```

where *Db2_installation_location* is the base installation directory for Db2 on your computer. For example, /opt/ibm/db2/V11.5.

Related information

- [db2licm - License management tool command](#)

Starting the installation programs

There are various methods to start the IBM Spectrum® Control installation programs and on various operating systems.

Installation program methods

You can use these methods to start the installation programs:

Installation images

You must select an installation image and download and extract these images to a location with adequate disk space.

Storage Resource agent

Contains the files to run local Storage Resource agent installations. You must download the file for the system on which you want the agent to be in.

- The Storage Resource agent location and operating system file name is the same as the product electronic image.
- The Storage Resource agent image is in *installation_dir/data/sra/operating_system_name*.

The operating systems that are supported for the Storage Resource agents are listed in the following table.

Table 1. Operating system for Storage Resource agents.

Operating system	Operating system name
Windows	windows
AIX®	aix_power
Linux® x86	linux_ix86
Linux for Power Systems Servers	linux_power
Linux s390	linux_s390

Important:

- If you are using IBM Spectrum Control electronic installation images that are decompressed in a directory that has spaces in the name, IBM Spectrum Control does not install correctly.

For example, if you have the following directory name:

- On Windows operating systems:

```
C:\Spectrum Control\disk\SC
```

- On AIX and Linux operating systems:

```
/temp/Spectrum Control/disk/SC
```

IBM Spectrum Control does not install correctly.

When you remove the spaces, and rename the directory, for example:

- On Windows operating systems:

```
C:\SpectrumControl\disk\SC
```

- On AIX and Linux operating systems:

```
/temp/SpectrumControl/disk/SC
```

IBM Spectrum Control installs correctly.

- If you are using IBM Spectrum Control electronic installation images for Linux or AIX operating systems and you download the images to a directory, ensure that your folder name *does not* contain a . at the end of the folder name.

For example, if you have a directory name

```
C:\SpectrumControl\disk\SC.
```

When you rename the folder and remove the . at the end of the name, for example:

```
C:\SpectrumControl\disk\SC.August3
```

IBM Spectrum Control installs correctly.

Starting the installation program on the Windows operating system

To start the IBM Spectrum Control installation program by using an electronic image, complete the following steps.

1. Download the image into a directory.
2. Extract the image files.
3. At the command prompt, enter `cd source_installation_directory\SC` and then enter `setup.bat`.

Starting the installation program on the Linux operating system

To start the IBM Spectrum Control installation program from the electronic image, complete the following steps.

1. Enter the following to create a directory:

```
mkdir /SpectrumControl
```

2. Download the image into the SpectrumControl directory.
3. Enter this command to extract the image files:

```
tar -xvzf filename.tar.gz
```

4. Ensure that the user name that you plan to use as the IBM Spectrum Control common user is in the `root` group and the `db2iadm1` group. If the user name is not in a group, you can add the user as shown in the following command example. In this example, the user name is `db2inst1`.

```
usermod -a -G root db2inst1
```

5. Navigate to the `/SpectrumControl/SC` directory by entering:

```
cd /SpectrumControl/SC
```

6. Set up your shell environment to point to the instance where the database repository is installed. Source the `db2profile` for the instance that you want. For example, if the Db2® instance is `db2inst1`, you can source the `db2profile` by entering:

```
./home/db2inst1/sqllib/db2profile
```

Remember: There is a space between . and /home.

7. Start the installation program by running the `./setup.bin` command from the same command shell that you used in [step 6](#).

Starting the installation program on the AIX operating system

To start the IBM Spectrum Control installation program from the electronic image, complete the following steps:

1. Enter the following to create a directory:

```
mkdir /SpectrumControl
```

2. Download the image into the SpectrumControl directory.
3. Enter this command to extract the image files:

```
tar -xvzf filename.tar.gz
```

4. Ensure that the user name that you plan to use as the IBM Spectrum Control common user is in the `system` group and the `db2iadm1` group. If the user name is not in a group, you can add the user as shown in the following command example. In this example, the user name is `db2inst1`.

```
usermod -G system db2inst1
```

Tip: If the user is already a member of secondary groups, the `usermod` command removes the `db2inst1` user name from all of the secondary groups except for the system group. To preserve the existing secondary group memberships, provide a comma separated list of secondary groups after the -G flag. For more information, enter `man usermod` at the command prompt.

5. Navigate to the `/SpectrumControl/SC` directory.

```
cd /SpectrumControl/SC
```

6. Set up your shell environment to point to the instance where the database repository is installed. Source the `db2profile` for the instance that you want. For example, if the Db2 instance is `db2inst1`, you can source the `db2profile` by entering:

```
. /home/db2inst1/sqllib/db2profile
```

Remember: There is a space between `.` and `/home`.

7. Start the installation program by running the `./setup.bin` command from the same command shell that you used in [step 6](#).

Installing IBM Spectrum Control in a single-server environment

Install IBM Spectrum® Control by using the installation program or the command line in silent mode.

Installation considerations

Before you begin the installation of IBM Spectrum Control, ensure that Db2® is installed. To learn how to install Db2, see [Installing Db2](#). If you install Db2 on a Windows server, reboot the system before you install IBM Spectrum Control.

When you install IBM Spectrum Control, a database that is called *TPCDB* is created. In this repository, the IBM Spectrum Control database schema is created, which describes the structure of the database repository. If you already have the *TPCDB* database on your system, IBM Spectrum Control assigns a new default database repository name with a number as a suffix. For example, *TPCDB1*.

If the operating system on which you are installing IBM Spectrum Control is sensitive to spaces in directory paths, change the default installation location to not include spaces, for example, `C:\apps\IBM\SC`.

- [Installing IBM Spectrum Control in a single-server Windows environment](#)
You can install IBM Spectrum Control by using the IBM Spectrum Control installation program.
- [Installing IBM Spectrum Control in a single-server AIX or Linux environment](#)
You can install IBM Spectrum Control as a root user or non-root user on an AIX or Linux operating system by using the IBM Spectrum Control installation program. To install IBM Spectrum Control as a non-root user, the user must have sudo privileges on the target AIX or Linux server.
- [Installing IBM Spectrum Control in a single-server environment by using silent mode](#)
You can install IBM Spectrum Control by using silent mode. The installation program does not display any configuration options during the installation process.

Installing IBM Spectrum Control in a single-server Windows environment

You can install IBM Spectrum® Control by using the IBM Spectrum Control installation program.

Procedure

To install IBM Spectrum Control in a single-server environment, complete the following steps:

1. Log on to the target server a user who as administrator privileges.
Note: Confirm that your Windows firewalls are disabled in the Control Panel.
2. When you install IBM Spectrum Control, the TPCDB database is created in the location that you specify in the Db2® `DFTDBPATH` variable.
To determine the current `DFTDBPATH` value, run the following command in a Db2 CLI window:

```
get dbm cfg
```

If you want to install this database in a location that is other than the default location, run the following command in a Db2 CLI window:

```
update dbm cfg using DFTDBPATH path
```

where *path* is the new value for the `DFTDBPATH` variable.

3. Use an extraction tool to extract the installation image files on the server.
4. In a command window, change the directory to the location of the IBM Spectrum Control installation files and enter the following command:

```
setup.bat
```

5. Follow the prompts in the installation program to install IBM Spectrum Control.

Results

If an error occurred during the IBM Spectrum Control installation process, review the IBM Spectrum Control installation log files to find details about the error.

Installing IBM Spectrum Control in a single-server AIX or Linux environment

You can install IBM Spectrum® Control as a root user or non-root user on an AIX® or Linux® operating system by using the IBM Spectrum Control installation program. To install IBM Spectrum Control as a non-root user, the user must have sudo privileges on the target AIX or Linux server.

Before you begin

Before you install IBM Spectrum Control, check the IBM Spectrum Control support site at <https://www.ibm.com/support/pages/node/6249361#Server> for the latest operating system support.

Ensure that X Window System is installed on the target AIX or Linux server. The X Window System is required to install IBM Spectrum Control by using the installation program.

If you want to install IBM Spectrum Control as a non-root user, have your system administrator complete the following tasks on the target server before you begin the installation:

- If you are installing IBM Spectrum Control on an AIX server, install sudo on the server. The default AIX distribution does not include sudo. To get the sudo installation package, go to [AIX Toolbox for Linux Applications](#). Find `sudo-version` in the Package column, and click RPM to download the package.
- Ensure that the non-root user has sudo privileges for the following commands on the server:
 - **xauth** (Enables the non-root user to provide access to their X Window System display for the root user.)
 - **db2ln** (Enables the non-root user to create the links for the Db2® system files.)
 - **setup.bin** (Enables the non-root user to launch the IBM Spectrum Control installation program.)

You must enter the full path to the commands when you configure the sudo privileges. For example:

```
username      ALL=(root)      SETENV:
/usr/bin/xauth, /opt/ibm/db2/V11.5/cfg/db2ln, /home/username/Downloads/IBMSC/SC/setup.bin
```

- If the sudo environment on the server uses the `secure_path` option, ensure that the following directories are included in the `secure_path` value for the non-root user who is installing IBM Spectrum Control:
 - `DB2_instance_owner_home_directory/sqllib/bin` (Example: `/home/db2inst1/sqllib/bin`)
 - `DB2_instance_owner_home_directory/sqllib/adm` (Example: `/home/db2inst1/sqllib/adm`)
 - `DB2_instance_owner_home_directory/sqllib/misc` (Example: `/home/db2inst1/sqllib/misc`)
- Ensure that the following requirements are met to enable the installation of IBM Spectrum Control by using the installation program:
 - The non-root user is able to use the X Window System on the target server and the X Window System **DISPLAY** environment variable is set correctly for the not-root user.
 - The **DISPLAY** environment variable is preserved in the sudo environment.
 - The root user has access to the X Window System display that is owned by the non-root user.These display requirements do not apply if you install IBM Spectrum Control in silent mode.
- If you are installing IBM Spectrum Control on an AIX server, ensure that the **ODMDIR** environment variable is preserved in the sudo environment.

Procedure

To install IBM Spectrum Control in a single-server environment, complete the following steps:

1. Log on to the target server as the root user or as a non-root user who has sudo privileges.
2. If you logged on as a non-root user, run the following command in a command shell:

```
sudo DB2_installation_directory/cfg/db2ln
```

For example:

```
sudo /opt/ibm/db2/V11.5/cfg/db2ln
```

3. When you install IBM Spectrum Control, the TPCDB database is created in the location that you specify in the Db2 **DFTDBPATH** variable. To determine the current **DFTDBPATH** value, run the following command in a Db2 CLI window:

```
get dbm cfg
```

If you want to install this database in a location that is other than the default location:

- a. Run the following command in a Db2 CLI window:

```
update dbm cfg using DFTDBPATH path
```

where *path* is the new value for the **DFTDBPATH** variable.

- b. Run the **db2stop** and **db2start** commands to stop and restart Db2.

4. Use an extraction tool to extract the installation image files on the server.
5. In a command shell, source the Db2 user profile **db2profile** for the Db2 instance owner. For example:

```
. /home/db2inst1/sqllib/db2profile
```

6. In a command shell, change the directory to the location of the IBM Spectrum Control installation program **setup.bin**, and then enter the following command:

Root user

```
./setup.bin
```

Non-root user

```
sudo -E ./setup.bin
```

7. Follow the prompts in the installation program to install IBM Spectrum Control.

Results

If an error occurred during the IBM Spectrum Control installation process, review the IBM Spectrum Control installation log files to find details about the error.

Installing IBM Spectrum Control in a single-server environment by using silent mode

You can install IBM Spectrum® Control by using silent mode. The installation program does not display any configuration options during the installation process.

Before you begin

If you plan to install IBM Spectrum Control as a non-root user, have your system administrator complete the tasks that are described in [Installing IBM Spectrum Control in a single-server AIX or Linux environment](#) on the target server before you begin the installation.

If you are installing on a Windows operating system, confirm that your Windows firewalls are disabled in the Control panel.

Procedure

To install IBM Spectrum Control in a single-server environment by using silent mode, complete the following steps:

1. Log on to the server where you want to install IBM Spectrum Control.
If you are installing on a Windows operating system, log on as a user who has administrator privileges. If you are installing on an AIX or Linux operating system, log on as the root user or as a non-root user who has *sudo* privileges.

2. If you logged on as an AIX or Linux non-root user, run the following command in a command shell:

```
sudo Db2_installation_directory/cfg/db2ln
```

For example:

```
sudo /opt/ibm/db2/V11.5/cfg/db2ln
```

3. When you install IBM Spectrum Control, the TPCDB database is created in the location that you specify in the Db2® **DFTDBPATH** variable.
To determine the current **DFTDBPATH** value, run the following command from a Db2 CLI window:

```
get dbm cfg
```

If you want to install this database in a location that is other than the default location:

- a. Run the following command in a Db2 CLI window:

```
update dbm cfg using DFTDBPATH path
```

where *path* is the new value for the **DFTDBPATH** variable.

- b. If you are installing on an AIX or Linux operating system, run the **db2stop** and **db2start** commands to stop and restart Db2.

4. Edit and save the appropriate response file.

For more information about editing the response file, see [Editing the response file](#).

5. If you are installing on an AIX or Linux operating system, in a command shell, source the Db2 user profile **db2profile** for the Db2 instance owner.
For example:

```
. /home/db2inst1/sqllib/db2profile
```

6. Run the silent mode installation program.

- For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```

where *language* can be one of the following values:

- Czech - cs
- English - en
- French - fr
- German - de
- Hungarian - hu
- Italian - it
- Japanese - ja
- Korean - ko
- Polish - pl
- Brazilian Portuguese - pt_BR
- Russian - ru
- Spanish - es
- Chinese (Simplified) - zh_CN
- Chinese (Traditional) - zh_TW

absolute_path_to_response_file is the absolute path to the response file. For example, the following command specifies the language and the path:

```
setup.bat -l fr -i silent -f c:\TPC\silent_SingleServerTypical.properties
```

- For AIX or Linux operating systems, run the following command:

Root user

```
./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

For example, the following command specifies the language and the path:

```
./setup.bin -l de -i silent -f /TPC/silent_SingleServerTypical.properties
```

Non-root user

```
sudo -E ./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

For example, the following command specifies the language and the path:

```
sudo -E ./setup.bin -l de -i silent -f /TPC/silent_SingleServerTypical.properties
```

If you do not specify a language or if the language is not specified correctly, the installation program uses the English language as the default.

7. Optional: Monitor the progress of the installation.

- To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
installation_dir\logs\traceTPCInstall.log
```

- To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

`installation_dir/logs/traceTPCInstall.log`

- **Editing the response file**

You must edit and save the appropriate response file when you install IBM Spectrum Control by using silent mode. The silent mode installation option for the response file is `-f absolute_path_to_response_file`. For example, in the Windows operating system, you enter `-f C:\installimage\silent_SingleServerTypical.properties`.

Editing the response file

You must edit and save the appropriate response file when you install IBM Spectrum® Control by using silent mode. The silent mode installation option for the response file is `-f absolute_path_to_response_file`. For example, in the Windows operating system, you enter `-f C:\installimage\silent_SingleServerTypical.properties`.

Use the following response files during a silent mode installation:

`silent_SingleServerTypical.properties`

This file specifies that all the IBM Spectrum Control components are installed on one server. You can customize this file by changing the ports for the Data server, Device server, Alert server, Export server, web server, and Storage Resource agent.

`silent_SingleServerCustom.properties`

This file specifies that all the IBM Spectrum Control components are installed on one server with the following customization options:

- Change the user name and password to install the database repository.
- Change the name of the database repository.
- Change the paths for the database repository.
- Change the path for the database repository log.
- Change the ports for the Data server, Device server, Alert server, Export server, web server, and Storage Resource agent.

`silent_MultipleServer.properties`

This file specifies the information that is required for a multiple-server environment. You can install the database repository on a remote server and install the other IBM Spectrum Control components on a local server.

Common response file parameters

These installation parameters are valid for all the response files.

`CHOSEN_INSTALL_TYPE="option"`

Specifies the installation type. Not all options are available in all the response files. [Table 1](#) shows you the options that you can specify.

Table 1. Options for the `CHOSEN_INSTALL_TYPE` parameter

Properties file	Valid options
<code>silent_SingleServerTypical</code>	<ul style="list-style-type: none"> • Single Server Install • License Upgrade
<code>silent_SingleServerCustom</code>	<ul style="list-style-type: none"> • Single Server Install • License Upgrade
<code>silent_MultipleServer</code>	<ul style="list-style-type: none"> • Multiple Server Install • License Upgrade

For example:

`CHOSEN_INSTALL_TYPE="Single Server Install"`

`LICENSE_ACCEPTED=false`

Specifies whether the user accepts the IBM Spectrum Control license agreement. The options are as follows:

`true`

The user accepts all the terms and conditions of the IBM Spectrum Control license agreement.

`false` or any other value

The user does not accept the IBM Spectrum Control license agreement. The installation program exits.

`USER_INSTALL_DIR=option`

Specifies the location where you want to install IBM Spectrum Control.

On Windows operating systems, you must use double backslashes. For example:

`USER_INSTALL_DIR=C:\\Program Files\\IBM\\TPC`

The default value for Windows operating systems is `C:\\Program Files\\IBM\\TPC`.

Note: If the operating system on which you are installing IBM Spectrum Control is sensitive to spaces in directory paths, change the default installation location to not include spaces, for example, `C:\\apps\\IBM\\SC`.

The default value for AIX® or Linux® operating systems is `/opt/IBM/TPC`.

`varSrvName=option`

Specifies the fully qualified host name of the server.

If your system is configured for dual stack networking (with both IPv4 and IPv6 addresses), IBM Spectrum Control defaults to IPv4 addressing.

Ensure that DNS is configured correctly on your server and verify that these files exist on your target server:

- For Windows operating systems, `C:\\Windows\\system32\\drivers\\etc\\hosts`.

- For AIX or Linux operating systems, /etc/hosts.

These files must have an entry similar to the following example:

```
myserver.example.mycompany.com myserver
```

varTPCPortRangeSP=port

Specifies the first port in a range of ports. The default beginning port is 9549. These ports are used by the IBM Spectrum Control servers, and the Storage Resource agent.

varCommonUsrID=user_name

Specifies the user name to install IBM Spectrum Control. The default user name for Windows operating system is db2admin. The default user name for AIX or Linux operating system is db2inst1.

The user name that you select must have the following operating system privileges:

- For Windows operating systems, the user name must be in the administrators group.
- For Linux operating systems, the user name must be in the root group.
- For AIX operating systems, the user name must be in the system group.

If you are using the common user name to also install the database repository, the user name must have these Db2® privileges:

- For Windows operating systems, the user name must be in the DB2ADMNS group.
- For AIX or Linux operating systems, the user name must be in the db2iadm1 group.

varCommonUsrPw=password

Specifies the password for the common user name.

varFullRollback=value

Specifies whether a full rollback or partial rollback is done when an installation failure occurs. These are the values:

- 0
Specifies that a partial rollback occurs. This value is the default.
- 1
Specifies that a full rollback occurs.

varUseLicenseKeyOnImage=value

Specifies whether to use the license key file that is present on the installation image. The value can be:

- 0
Uses the license key that is present on the installation image (default value)
- 1
Provides the location of the license key file

Depending on the license package you purchased, one of these license files is included with IBM Spectrum Control in the product electronic image:

```
nodeLock.AE
    IBM Spectrum Control license
nodeLock.AS
    IBM Spectrum Control Select Edition license
```

varLicenseKeyFile=location_of_license_key_file

Specifies the path for the license key file. This value must be set if the varUseLicenseKeyOnImage=1 parameter is specified.

On the Windows operating system, you must use double backslashes. For example:

```
C:\\license\\key\\nodeLock.AE
```

Response file parameters for advanced customization of the database repository

These parameters are used for the **silent_SingleServerCustom.properties** and **silent_MultipleServer.properties** response files to configure the database repository.

varDBAdmUsr=user_name

Specifies the user name that you must use to install the IBM Spectrum Control database. The user who logs in with this user name must be the owner of the Db2 instance where the database is created.

This user name must have these Db2 privileges:

- For Windows operating systems, the user name must be in the DB2ADMNS group. The default user name is db2admin.
- For AIX or Linux operating systems, the user name must be in the db2iadm1 group. The default user name is db2inst1.

varDBAdmPW=password

Specifies the password that is associated with the Db2 user name.

varDBName=database_name

Specifies the name of the IBM Spectrum Control database. The default database name is TPCDB.

Specify the database name by using the following rules:

- The name can contain the following characters:
 - a - z
 - A - Z
 - 0 - 9
- The name must be 1 - 8 characters long.

Restriction: The database name cannot have the following conditions:

- Contain a space or a blank.
- Begin with SYS, DBM, or IBM®.

- Be the name of an existing database.

`varDBPath=database_location, database_location, ...`
Specifies the absolute paths for the database.

The database path must follow these rules:

- The path cannot be longer than 242 bytes.
- The path must be an absolute path and not a relative path, for example, C:\

You can specify a maximum of 10 database paths and separate them with commas. For example, in the Windows operating system, you would specify `C: ,D: .` In the AIX or Linux operating system, you would specify `/home/db2inst1,/testPath`.

- For example, on the Windows operating system, if the database is on the C drive, you can add a D drive, an E drive. On Windows operating systems, you must use double backslashes. For example:

```
C: \DB2 \TPCDB ,D: \TPCDB ,E: \TPCDB
```

- For example, on the AIX operating system:

```
/home/db2inst1,/testPath
```

`varDBLogPath=log_location`

Specifies the absolute path where the database log files are stored. The default value of the variable is set to `<dftdbpath>/<instance name>/<varDBName>/SQLLOG`.

For example:

- On Windows operating systems, `c:\DB2\TPCDB\SQLLOG`
- On AIX operating systems, `/home/db2inst1/db2inst1/TPCDB/SQLLOG`

Response file parameter for the multiple-server environment

This parameter is valid for the `silent_MultipleServer.properties` file.

`varMultipleServerComponentList=components`

Specifies the IBM Spectrum Control components to install.

`database`

Installs the IBM Spectrum Control database repository.

`servers`

Installs the Data server, Device server, Alert server, Web server, the Storage Resource agent, CLI, and the IBM Spectrum Control GUI.

Response file parameters for the remote database

The following parameters are used only in the `silent_MultipleServer.properties` response file to access the IBM Spectrum Control database on a remote server. Remember: The Db2 database must be installed before you can use these parameters to access the database.

`varRemoteDBSrvName=host_name`

Specifies the fully qualified host name or IP address of the remote server where the IBM Spectrum Control database repository is installed. IBM Spectrum Control accepts both IPv4 and IPv6 addresses. If you have a system that is configured for dual stack networking (with both IPv4 and IPv6 addresses), IBM Spectrum Control defaults to IPv4 addressing.

For example, an IPv6 address in long form can be `2001:DB8:0:0:0:0:0:0`. An IPv6 address in short form can be `2001:DB8::` or `2001:DB8:0:0:0:0:0:0`.

`varRemoteDBPort=port`

Specifies the Db2 port on the remote server. The default port is 25000. Port 50000 can also be used.

`varRemoteDBAdmUsr=user_name`

Specifies the administrative user name to connect to the remote database. The default user name for the Windows operating system is `db2admin`. The default user name for the AIX or Linux operating system is `db2inst1`.

`varRemoteDBAdmPW=password`

Specifies the password for the user name.

`varRemoteDBName=database_name`

Specifies the IBM Spectrum Control database on the remote server. The default name is `TPCDB`.

Related reference

- [Ports used by IBM Spectrum Control](#)

Installing IBM Spectrum Control in a multiple-server environment

You can install IBM Spectrum® Control by using the installation program or by using silent mode from the command line. Note, installing IBM Spectrum Control by using console mode is not supported.

A multiple-server environment is ideal if you are monitoring large storage environments, where one server is not sufficient to manage the IBM Spectrum Control components.

If you have different administrators for Db2® and IBM Spectrum Control, this environment allows the database repository and Db2 to be installed and managed on a separate server.

Installation considerations

When you install IBM Spectrum Control, a database that is called TPCDB is created. In this repository, the IBM Spectrum Control database schema is created, which describes the structure of the database repository. If you already have the TPCDB database on your system, IBM Spectrum Control assigns a new default database repository name with a number as a suffix. For example, TPCDB1.

If the operating system on which you are installing IBM Spectrum Control is sensitive to spaces in directory paths, change the default installation location to not include spaces, for example, C:\apps\IBM\SC.

- [Installing IBM Spectrum Control with a remote database by using the installation program](#)
You can install IBM Spectrum Control in a multiple-server environment by using the installation program.
- [Installing IBM Spectrum Control with a remote database by using silent mode](#)
You can install IBM Spectrum Control in a multiple-server environment by using silent mode.

Related tasks

- [Installing IBM Spectrum Control with a remote database by using the installation program](#)
- [Installing IBM Spectrum Control with a remote database by using silent mode](#)

Installing IBM Spectrum Control with a remote database by using the installation program

You can install IBM Spectrum® Control in a multiple-server environment by using the installation program.

Before you begin

For this procedure, the terms *Server A* and *Server B* denote the two servers. If you plan to install IBM Spectrum Control as a non-root user, have your system administrator complete the tasks that are described in [Installing IBM Spectrum Control in a single-server AIX or Linux environment](#) on both servers before you begin the installation.

About this task

Server A will have Db2® and the IBM Spectrum Control database repository installed. Server B will have the IBM Spectrum Control servers installed.

Procedure

To install IBM Spectrum Control with a remote database by using the installation program, complete the following steps:

1. Complete the following steps on Server A:
 - a. Log on to Server A.
If you are installing on a Windows operating system, log on as a user who has administrator privileges. If you are installing on an AIX or Linux operating system, log on as the root user or as a non-root user who has sudo privileges.
Note: Confirm that your Windows firewalls are disabled in the Control Panel.
 - b. Install Db2 on Server A.
For more information about installing Db2, see [Installing Db2](#). If you install Db2 on a Windows server, restart the server before you install IBM Spectrum Control.
 - c. If you logged on as an AIX or Linux non-root user, run the following command in a command shell:

```
sudo DB2_installation_directory/cfg/db2ln
```

For example:

```
sudo /opt/ibm/db2/V11.5/cfg/db2ln
```

- d. When you install IBM Spectrum Control, the TPCDB database is created in the location that you specify in the Db2 **DFTDBPATH** variable.
To determine the current **DFTDBPATH** value, run the following command in a Db2 CLI window:

```
get dbm cfg
```

If you want to install this database in a location that is other than the default location:

- i. Run the following command in a Db2 CLI window:

```
update dbm cfg using DFTDBPATH path
```

where *path* is the new value for the **DFTDBPATH** variable.

- ii. If you are installing on an AIX or Linux operating system, run the **db2stop** and **db2start** commands to stop and restart Db2.

- e. If you are installing on an AIX or Linux operating system, in a command shell, source the Db2 user profile **db2profile** for the Db2 instance owner.
For example:

```
. /home/db2inst1/sqllib/db2profile
```

- f. Start the IBM Spectrum Control installation program on Server A.

In a command window or command shell, change the directory to the location of the IBM Spectrum Control installation files and enter the following command:

Windows

```
setup.bat
```

AIX or Linux root user

```
./setup.bin
```

AIX or Linux non-root user

```
sudo -E ./setup.bin
```

- g. On the Before You Begin page, because only the IBM Spectrum Control database repository will be installed on Server A, click Next.
 - h. On the Choose Installation Location and Type page, select Multiple servers and click Next.
 - i. On the Multiple Server Option: Select the Components page, select Database repository and click Next.
 - j. Follow the prompts in the installation program to install the IBM Spectrum Control database repository on Server A.
 - k. After the installation is finished, review the message log on Server A to ensure that no errors occurred.
2. Complete the following steps on Server B:
- a. Log on to Server B.
If you are installing on a Windows operating system, log on as a user who has administrator privileges. If you are installing on an AIX or Linux operating system, log on as the root user or as a non-root user who has sudo privileges.
 - b. Start the IBM Spectrum Control installation program on Server B.
In a command window or command shell, change the directory to the location of the IBM Spectrum Control installation files and enter the following command:
Windows

```
setup.bat
```


AIX or Linux root user

```
./setup.bin
```


AIX or Linux non-root user

```
sudo -E ./setup.bin
```
 - c. On the Choose Installation Location and Type page, select Multiple servers and click Next.
 - d. On the Multiple Server Options: Select the Components page, select IBM Spectrum Control Servers and click Next.
Attention: **IBM Spectrum Control servers** include the following components:
 - IBM Spectrum Control GUI
 - CLI
 - Storage Resource agent
 - e. Follow the prompts in the installation program to install the IBM Spectrum Control Servers on Server B.
 - f. After the installation program is finished, review the message log on Server B to ensure that no errors occurred.

Related reference

- [Planning for IBM Spectrum Control authentication and authorization](#)
- [Installing IBM Spectrum Control in a multiple-server environment](#)
- [Starting the installation programs](#)
- [IBM Spectrum Control components](#)
- [Verifying the installation](#)
- [Reviewing the log files to resolve installation issues](#)

Installing IBM Spectrum Control with a remote database by using silent mode

You can install IBM Spectrum® Control in a multiple-server environment by using silent mode.

Before you begin

For this procedure, the terms *Server A* and *Server B* denote the two servers. If you plan to install IBM Spectrum Control as a non-root user, have your system administrator complete the tasks that are described in [Installing IBM Spectrum Control in a single-server AIX or Linux environment](#) on both servers before you begin the installation.

If you are installing IBM Spectrum Control in silent mode on a Windows operating system, confirm that your Windows firewalls are disabled in the Control Panel.

About this task

Server A will have Db2® and the IBM Spectrum Control database repository installed. Server B will have the IBM Spectrum Control servers installed.

To install IBM Spectrum Control in a multiple-server environment by using silent mode, complete the following steps:

Procedure

1. Complete the following steps on Server A:
 - a. Log on to Server A.
If you are installing on a Windows operating system, log on as a user who has administrator privileges. If you are installing on an AIX or Linux operating system, log on as the root user or as a non-root user who has sudo privileges.
 - b. Install Db2 on Server A.
For more information about installing Db2, see [Installing Db2](#). If you install Db2 on a Windows server, make sure that you restart the server before you install IBM Spectrum Control.
 - c. If you logged on as an AIX or Linux non-root user, run the following command in a command shell:

```
sudo DB2_installation_directory/cfg/db2ln
```

For example:

```
sudo /opt/ibm/db2/V11.5/cfg/db2ln
```

- d. When you install IBM Spectrum Control, the TPCDB database is created in the location that you specify in the Db2 **DFTDBPATH** variable. To determine the current **DFTDBPATH** value, run the following command in a Db2 CLI window:

```
get dbm cfg
```

If you want to install this database in a location that is other than the default location:

- i. Run the following command in a Db2 CLI window:

```
update dbm cfg using DFTDBPATH path
```

where *path* is the new value for the **DFTDBPATH** variable.

- ii. If you are installing on an AIX or Linux operating system, run the **db2stop** and **db2start** commands to stop and restart Db2.

- e. On Server A, set the following parameters in the **silent_MultipleServer.properties** response file to install the remote database repository.

- LICENSE_ACCEPTED=true
- CHOSEN_INSTALL_TYPE="Multiple Server Install"
- varMultipleServerComponentList=database
- USER_INSTALL_DIR=option
- varSrvName=option
- varDBAdmUsr=user_name
- varDBAdmPW=password
- varDBName=database_name
- varDBPath=database_location
- varDBLogPath=log_location

For more information about editing the **silent_MultipleServer.properties** file, see [Editing the response file](#).

- f. Save the response file.

- g. If you are installing on an AIX or Linux operating system, in a command shell, source the Db2 user profile **db2profile** for the Db2 instance owner. For example:

```
. /home/db2inst1/sqllib/db2profile
```

- h. Run the silent mode installation program on Server A.

- For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```

where *language* can be one of the following values:

- Czech - cs
- English - en
- French - fr
- German - de
- Hungarian - hu
- Italian - it
- Japanese - ja
- Korean - ko
- Polish - pl
- Brazilian Portuguese - pt_BR
- Russian - ru
- Spanish - es
- Chinese (Simplified) - zh_CN
- Chinese (Traditional) - zh_TW

absolute_path_to_response_file is the absolute path to the response file. For example, the following command specifies the language and the path:

```
setup.bat -l fr -i silent -f c:\TPC\silent_MultipleServer.properties
```

- For AIX or Linux operating systems, run the following command:

Root user

```
./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

For example, the following command specifies the language and the path:

```
./setup.bin -l de -i silent -f /TPC/silent_MultipleServer.properties
```

Non-root user

```
sudo -E ./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

For example, the following command specifies the language and the path:

```
sudo -E ./setup.bin -l de -i silent -f /TPC/silent_MultipleServer.properties
```

If you do not specify a language or if the language is not specified correctly, the installation program uses the English language as the default.

- i. Optional: Monitor the progress of the installation.

- To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
installation_dir\logs\traceTPCInstall.log
```

- To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

```
installation_dir/logs/traceTPCInstall.log
```

Tip: If there are preinstallation errors, you can review the **lax*out.txt** and **lax*err.txt** files that are in the **/tmp** and **%TEMP%** directories.

2. Complete the following steps on Server B:

a. Log on to Server B.

If you are installing on a Windows operating system, log on as a user who has administrator privileges. If you are installing on an AIX or Linux operating system, log on as the root user or as a non-root user who has sudo privileges.

b. On Server B, set the following parameters in the `silent_MultipleServer.properties` response file:

- `LICENSE_ACCEPTED=true`
- `CHOSEN_INSTALL_TYPE="Multiple Server Install"`
- `varMultipleServerComponentList=servers`
- `USER_INSTALL_DIR=option`
- `varSrvName=option`
- `varTPCPortRangeSP=port`
- `varCommonUsrID=user_name`
- `varCommonUsrPw=password`
- `varRemoteDBSrvName=host_name`
- `varRemoteDBPort=port`
- `varRemoteDBAdmUsr=user_name`
- `varRemoteDBAdmPW=password`
- `varRemoteDBName=database_name`

For more information about editing the `silent_MultipleServer.properties` file, see [Editing the response file](#).

c. Run the silent mode installation program on Server B.

- For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```

where *language* can be one of the following values:

- Czech - cs
- English - en
- French - fr
- German - de
- Hungarian - hu
- Italian - it
- Japanese - ja
- Korean - ko
- Polish - pl
- Brazilian Portuguese - pt_BR
- Russian - ru
- Spanish - es
- Chinese (Simplified) - zh_CN
- Chinese (Traditional) - zh_TW

`absolute_path_to_response_file` is the absolute path to the response file. For example, the following command specifies the language and the path:

```
setup.bat -l fr -i silent -f c:\TPC\silent_MultipleServer.properties
```

- For AIX or Linux operating systems, run the following command:

Root user

```
./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

For example, the following command specifies the language and the path:

```
./setup.bin -l de -i silent -f /TPC/silent_MultipleServer.properties
```

Non-root user

```
sudo -E ./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

For example, the following command specifies the language and the path:

```
sudo -E ./setup.bin -l de -i silent -f /TPC/silent_MultipleServer.properties
```

If you do not specify a language or if the language is not specified correctly, the installation program uses the English language as the default.

d. Optional: Monitor the progress of the installation.

- To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
installation_dir\logs\traceTPCInstall.log
```

- To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

```
installation_dir/logs/traceTPCInstall.log
```

Note: If there are preinstallation errors, you can review the `lax*out.txt` and `lax*err.txt` files that are in the `/tmp` and `%TEMP%` directories.

Related reference

- [Planning for IBM Spectrum Control authentication and authorization](#)
- [Installing IBM Spectrum Control in a multiple-server environment](#)
- [Starting the installation programs](#)
- [IBM Spectrum Control components](#)
- [Verifying the installation](#)
- [Reviewing the log files to resolve installation issues](#)

Accessing IBM Spectrum Control server from a remote computer using Command Line Interface

Accessing client components on a remote computer helps you complete storage management tasks from a computer that does not have a IBM Spectrum® Control server that is installed on it.

Procedure

To access IBM Spectrum Control server on a remote computer, complete the following steps:

1. Go to the `client_images` subdirectory in the IBM Spectrum Control installation directory.

- On AIX® or Linux® operating systems, here is an example of a subdirectory:

```
/opt/IBM/TPC/client_images
```

- On Windows operating systems, here is an example of a subdirectory:

```
C:\Program Files\IBM\TPC\client_images
```

2. To create a client image, run one of the following commands for the operating system of the computer on which you want to access the command-line interface (CLI):

- On AIX or Linux operating systems, run the following command:

```
createClientImage.sh [aix | linux | windows | all]
```

When you include the `[aix | linux | windows | all]` parameters, you can create AIX, Linux, and Windows client images on a computer that is running AIX or Linux operating systems.

- On Windows operating systems run the following command:

```
createClientImage.bat [aix | linux | windows | all]
```

When you include the `[aix | linux | windows | all]` parameters, you can create AIX, Linux, and Windows client images on a computer that is running AIX and Linux or Windows operating systems.

Tip: Each parameter creates one of the following compressed client image files:

`aix`

Creates a `TPC_CLIENT_AIX.tar` image.

`linux`

Creates a `TPC_CLIENT_LINUX.zip` image.

`windows`

Creates a `TPC_CLIENT_WIN.zip` image.

`all`

Creates all of the client images.

3. Extract the appropriate compressed file on the remote computer.

- On AIX operating systems, extract the `TPC_CLIENT_AIX.tar` file to a folder, for example:

```
/opt/IBM/TPCclient
```

- On Linux operating systems, extract the `TPC_CLIENT_LINUX.zip` file to a folder, for example:

```
/opt/IBM/TPCclient
```

- On Windows operating systems, extract the `TPC_CLIENT_WIN.zip` file to a folder, for example:

```
c:/Program Files/IBM/TPCclient
```

4. Ensure that the following directories appear in the folder where you extracted the client image compressed file:

- `cli`
- `jre`

The `cli` folder contains the files to start the IBM Spectrum Control command-line interface and the Replication command-line interface. The `jre` folder contains other files that are necessary to start the command-line interface.

5. To start the IBM Spectrum Control command-line interface, run one of the following commands:

- On AIX or Linux operating systems:

```
/opt/IBM/TPCclient/cli/tpctool.sh
```

- On Windows operating systems:

```
C:\Program Files\IBM\TPCclient\cli\tpctool.bat
```

Installing IBM Spectrum Control on a Windows domain

You can install IBM Spectrum® Control and Db2® by using a Windows domain or a local user account.

Installing Db2 and IBM Spectrum Control by using local user accounts

Prerequisite: If the local user account is used to install the database repository, this account must be a member of the Db2 administrator (DB2ADMNS) group. To install IBM Spectrum Control and Db2 by using local user account, complete the following steps:

1. Install Db2 by using a local user account.
2. Log on to the Windows domain computer on which you plan to install IBM Spectrum Control by using a local user account that is a member of the local Administrators group and the local Db2 Administrators (DB2ADMNS) group.
3. On the Single Server Installation Information page, in the User name field, enter a local user account that is a member of local administrator group.

Installing Db2 by using a local user account and IBM Spectrum Control by using a domain user account

Prerequisite: If the local user account is used to install the database repository, this account must be a member of the Db2 administrator (DB2ADMNS) group. To install Db2 with a local user account and IBM Spectrum Control with a domain user account, complete the following steps:

1. Install Db2 by using a local user account.
2. Log on to the Windows domain computer on which you plan to install IBM Spectrum Control by using a domain user account that is a member of the local Administrators group and the local Db2 Administrators (DB2ADMNS) group.
This user must have the Act as part of the operating system and Login as a service permissions set in the security policy.
3. On the Single Server Installation Information page, in the User name field, enter a domain user account (with a Windows domain name prefix) that is a member of the local administrators group.
4. Click Configure Database repository.
5. On the Configure the Database Repository page, enter the Db2 user name **without** a Windows domain name prefix, enter the password for the Db2 user, and click Validate.
6. When the validation is complete, click OK to return to the Single Server Installation Information page.

Restriction: If the domain security policy overrides the local security policy, you must set the indicators on the domain controller computer.

Installing Db2 and IBM Spectrum Control by using domain user accounts

Prerequisite: If the domain user account is used to install the database repository, this user account must have Db2 SYSADM authority. To install IBM Spectrum Control and Db2 by using domain user accounts complete the following steps:

1. Install Db2 by using a domain user account.
2. Log on to the Windows domain computer on which you plan to install IBM Spectrum Control by using a domain user account that is a member of the local Administrators group and the domain Db2 Administrators (DB2ADMNS) group.
3. On the Single Server Installation Information page, in the User name field, enter a domain user account (with a Windows domain name prefix) that is a member of the local administrators group.
4. Click Configure Database repository.
5. On the Configure the Database Repository page, enter the Db2 user name **without** a Windows domain name prefix, enter the password for the Db2 user, and click Validate.
6. When the validation is complete, click OK to return to the Single Server Installation Information page.

Resolving Windows domain prevalidation errors

Adding Windows domain user accounts to a local group, such as Administrator or DB2ADMNS, and later deleting these user accounts from the domain controller computer can cause an error when the account is validated.

When you delete user accounts from the domain controller computer, a reference to the user account is created as a security ID on the computer that is a member of a Windows domain. If these references are not removed, the IBM Spectrum Control installation can fail.

To locate the references that have a security ID, complete these steps:

1. On the domain member computer, click Start > Control Panel > Administrative Tools > Computer Management. If you are prompted by a User Account Control window, click Yes.
2. In the Computer Management navigation tree view, expand the Local Users and Groups node and select Groups.
3. In the Name column, look for a security ID like the following example:

```
S-1-5-21-337177553-1671989427-887411491-500
```

4. To remove a reference, select the security ID and click Delete.
 - a. To add the user account back to the group, click Add. The user account now belongs to the group.

When a prevalidation error occurs, depending on the type of user account and the group, error messages might be displayed:

- BPCIN0109E An unexpected error occurred.....
For more information, review the log files and see the IBM Documentation.

To resolve this situation, use these steps:

1. On the domain controller computer, open Users and Groups on the Active Directory.
2. Delete all groups from the user account, and then add the groups again.
3. Log on to the computer again that is the member of the domain.

You have to add the DB2ADMNS rights to the user account again.

Locate the installation log files, lax*-out.txt and lax*-err.txt, for more details about the root cause of the error message. For example,

```
name=DB2ADMNS An error occurred while enumerating the groups.  
The group could be found.
```

- BPCIN0244E An error occurred while enumerating the local administrator group membership. On the current computer, on the Properties page of this group, remove the user names that are displayed with a security ID (SID). An example of a SID is S-1-5-21-337177553-1671989427-887411491-500 and is used instead of a user name.
- BPCIN0245E An error occurred while enumerating the local DB2 administrator group membership. On the current computer, on the Properties page of this group, remove the user names that are displayed with a security ID (SID). An example of a SID is S-1-5-21-337177553-1671989427-887411491-500 and is used instead of a user name.

For more information about these messages, go to IBM® Docs and search for the message number.
In the error message window, click OK to resume the installation program or Quit to restart the IBM Spectrum Control.

Related reference

- [Windows domain and local user accounts](#)
- [Adding a computer to the Windows domain](#)
- [Installing Db2 by using a Windows domain user account](#)
- [Creating a Windows domain common user account for IBM Spectrum Control](#)
- [Granting Db2 SYSADM authority to a Windows domain user account](#)

Verifying the connection to the domain controller computer by using the Dcdiag tool

Use the Dcdiag command line tool to help you determine whether the domain controller computer is registered with the domain name server (DNS), whether the controller can be pinged, and whether the controller has Lightweight Directory Access Protocol (LDAP) connectivity.

About this task

The Dcdiag tool is build into Windows Server versions 2019, 2016, and 2012 R2.

To use the Dcdiag tool, complete the following steps:

Procedure

1. On a computer that is a member of a Windows domain, go to [https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-r2-and-2012/cc731968\(v=ws.11\)](https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2012-r2-and-2012/cc731968(v=ws.11)).

To download the Dcdiag tool, complete the following steps:

- a. Go to the Windows Start menu and start the Server Manager.
- b. On the Server Manager Dashboard, click Add roles and features.
- c. On the Before You Begin page, click Next.
- d. On the Installation Type page, verify that Role-based or feature-based installation' is selected and click Next.
- e. On the Server Selection page, ensure that your server is selected and click Next.
- f. On the Server Roles page, click Next.
- g. On the Features page, go to Remote Server Administration Tools -> Role Administration Tools -> AD DS.
- h. Select AD LDS Tools -> AD DS Tools, and click AD DS Snap-Ins and then Command Line Tools. Click Next.
- i. On the Confirmation page, click Install.
- j. Verify that your installation is complete, then click Close.

2. In a command prompt, enter the following command:

```
DCDIAG /TEST:DNS /V /E /S:domaincontroller
```

Where *domaincontroller* is the name of the domain controller computer, for example, *TPCDomain*.

Important: You can run the command only if you are logged in with a Windows domain user name. If you try to run the command when you are logged in with a local user name, the following error message is displayed:

```
Connecting to directory service on server 10.10.10.10  
LDAP bind failed with error 1326,  
Logon failure: unknown user name or bad password.
```

Example

When you run the **DCDIAG** command, and you connect to the domain controller computer, the output contains text such as `passed test Connectivity`. This text is displayed even if text such as `system1.srm.tpc.example.com failed test DNS` is displayed at the end of the output.

The following output shows that the connection was successful:

```
Starting test: Connectivity  
* Active Directory LDAP Services Check  
*** Warning: could not confirm the identity  
of this server in the directory versus  
the names returned by DNS servers. If  
there are problems accessing this directory  
server then you may need to check that this  
server is correctly registered with DNS
```

Installing IBM Spectrum Control and associated products using minimal space on the Windows C: drive

In some cases, you might not want to install Db2® or IBM Spectrum® Control on the C: drive.

For example, you might have limited space on the drive. You can install these products on a different drive that meets the space requirements for the products. You can also minimize the space that is used on the C: drive during the installation process by directing temporary files to another drive.

The procedures that are provided apply to installing the products on a single server using the appropriate installation programs. The procedures can be modified to suit different installation scenarios.

Installing Db2

Complete the following steps to install Db2 with minimal space used on the C: drive:

1. Open a command window and complete the following steps in the same window session:
 - a. At the command prompt, change the directory to the location of the Db2 Setup program setup.exe file.
 - b. Run the following commands to change the TEMP and TMP environment variable settings to the drive and directory where you want to direct the temporary files that are created during the installation of Db2.
Important: The directory that is specified in these commands must be an existing directory. If the directory does not exist, you must create it before you start the installation of Db2.

```
set TEMP=drive_and_directory
set TMP=drive_and_directory
```

For example:

```
set TEMP=Q:\Temp
set TMP=Q:\Temp
```

2. Start the Db2 Setup program and follow the prompts in the program to complete a typical installation on a drive other than the C: drive. Ensure that you set the correct drive and directory on the following pages:
 - Select the installation folder
 - Select the IBM SSH server installation folder and startup option

After the installation, the files and directories that are related to Db2 use approximately 75 MB of space on the C: drive.

Installing IBM Spectrum Control

Complete the following steps to install IBM Spectrum Control with minimal space used on the C: drive:

1. Open a command window and complete the following steps in the same window session:
 - a. At the command prompt, change the directory to the location of the IBM Spectrum Control installation program setup.bat file.
 - b. Run the following commands to change the TEMP and TMP environment variable settings to the drive and directory where you want to direct the temporary files that are created during the installation of IBM Spectrum Control.
Important: The directory that is specified in these commands must be an existing directory. If the directory does not exist, you must create it before you start the installation of IBM Spectrum Control.

```
set TEMP=drive_and_directory
set TMP=drive_and_directory
```

For example:

```
set TEMP=Q:\Temp
set TMP=Q:\Temp
```

2. Start the IBM Spectrum Control program and follow the prompts to complete the installation on a drive other than the C: drive. Ensure that you set the correct drive and directory on the following pages:
 - Choose the Installation Location and Type page in the IBM Spectrum Control installation program.
 - Configure the Database Repository page in the IBM Spectrum Control installation program.

After the installation, the files and directories that are related to IBM Spectrum Control use approximately 5 MB on the C: drive.

Verifying the installation

After you install IBM Spectrum® Control, you can verify whether the installation was successful.

Use the following questions to guide you through the verification process.

Remember the following information when you are verifying the installation:

- You must use a user name with Db2® administrative privileges to run any Db2 commands.
- You must run any Db2 commands on the Db2 command line.

Was the TPCDB database created?

Run the following Db2 command:


```
db2 list db directory
```

This command returns information about the TPCDB database.

Are the tables and views present in the TPCDB database?

Run the following Db2 commands:

```
db2 connect to TPCDB
db2 list tables for schema TPC
db2 disconnect TPCDB
```

These commands indicate that you can connect to and disconnect from the IBM Spectrum Control database repository and lists the tables and views in the TPCDB database.

Is the Data server running?

On the Windows operating system:

Go to the Services listing and find IBM Spectrum Control - Data Server.

On the AIX® or Linux® operating system:

Run the following command:

```
ps -eaf | grep DataServer.jar
```

This command returns a list of the Data server processes in the process table.

Is the Device server running?

On the Windows operating system:

1. Go to the Services listing and find IBM Spectrum Control - Device Server.
2. In a Windows command window, go to the following directory:

```
installation_dir\wlp\bin\
```

where *installation_dir* is the top-level directory in which you installed IBM Spectrum Control (for example, *C:\Program Files\IBM\TPC*).

3. Run the following command:

```
server.bat status deviceServer
```

This command returns a status of *running*.

On the AIX or Linux operating system:

1. In a command-line window, go to the following directory:

```
installation_dir/wlp/bin/
```

where *installation_dir* is the top-level directory in which you installed IBM Spectrum Control (for example, */opt/IBM/TPC/*).

2. Run the following command:

```
./server status deviceServer
```

This command returns a status of *running*.

It can take several minutes after the Device server starts for all services to be running.

Are all Device server services running?

Go to

```
https://localhost:port/ITSRM/ServiceManager
```

Where *port* is the value of the *WC_defaulthost_secure* key in the *installation_dir/device/portdef.props* file. The default port number is 9551.

Is the Storage Resource agent running?

On the Windows operating system:

Go to the Services window and find IBM Spectrum Control Storage Resource Agent.

On the AIX or Linux operating system:

Run the following command:

```
ps -eaf | grep Agent
```

This command returns a list of the Storage Resource agent processes in the process table.

Is the web server running?

On the Windows operating system:

Go to the Services window and find IBM Spectrum Control - Web Server.

On the AIX or Linux operating system:

Run the following command:

```
ps -eaf | grep webServer
```

This command returns a list of the web server processes in the process table.

Can I log on to the IBM Spectrum Control GUI?

Open a web browser and enter the following web address:

```
https://host_name:port/srm
```

Where *host_name* is the fully qualified host name of the computer where the IBM Spectrum Control Web server is installed and *port* is the value of the *httpsPort* key in the *installation_dir/wlp/usr/servers/webServer/server.xml* file. The default port is 9569.

Reviewing the log files to resolve installation issues

If an error occurs during IBM Spectrum® Control installation, you can review the error log files to resolve issues and continue the installation.

To resolve errors that occur during the preinstall steps in the IBM Spectrum Control installation program, complete the following steps:

1. Navigate to one of the following directories:
 - For the Windows operating system, go to %TEMP% directory
 - For the AIX® or Linux® operating system, go to /tmp directory
2. Review the following log files:
 - lax-xxxxx-out.txt
 - lax-xxxxx-err.txt

Tip: The files are InstallAnywhere log files and are automatically generated. The xxxxx are numeric characters that are assigned by the InstallAnywhere software. To resolve errors after the IBM Spectrum Control installation starts, complete the following steps:

1. Navigate to one of the following default directories:
 - For the Windows operating system, go to `C:\Program Files\IBM\TPC\logs`
 - For the AIX or Linux operating system, go to `/opt/IBM/TPC/logs`
2. Review these log files:
 - msgTPCInstall.log
 - traceTPCInstall.log
 - sra\install\agent.trace
 - sra\install\agent_number.log
3. For agent log files, navigate to one of the following default directories:
 - For the Windows operating system, go to `SRA_location\log`
 - For the AIX or Linux operating system, go to `SRA_location/log`where `SRA_location` is where the Storage Resource agent is installed.

If an installation error occurred and you decide to do a partial rollback, the msgTPCInstall.log and traceTPCInstall.log files are updated with additional information when you resume the installation.

Changing languages

You can change the operating system language and web browser that determines the language in which IBM Spectrum® Control services and the IBM Spectrum Control GUI are displayed.

For IBM Spectrum Control, the following languages are supported:

- English
 - Czech
 - French
 - German
 - Hungarian
 - Italian
 - Japanese
 - Korean
 - Polish
 - Brazilian Portuguese
 - Russian
 - Spanish
 - Chinese (Simplified)
 - Chinese (Traditional)
- [Changing the language of the IBM Spectrum Control GUI](#)
Change the language that is displayed in your web browser for the IBM Spectrum Control GUI.
 - [Changing the operating system language for Windows](#)
Change the operating system language for Windows.
 - [Changing the operating system language for AIX](#)
Change the operating system language for AIX®.

Changing the language of the IBM Spectrum Control GUI

Change the language that is displayed in your web browser for the IBM Spectrum® Control GUI.

Procedure

1. Choose one of these options to access the settings page in the web browser:

Option	Description
Mozilla Firefox (ESR 31 and later)	In the menu bar, click Tools > Options.
Google Chrome	Click the customization control in the upper-right corner of the page and click Settings.
Internet Explorer (v11 and later)	In the menu bar, click Tools > Internet Options.

2. Access the language settings in the web browser:

Option	Description
Mozilla Firefox (ESR 31 and later)	Scroll to the Language and Appearance section and click Choose.
Google Chrome	Scroll to the bottom of the Settings page and click Advanced.
Internet Explorer (v11 and later)	In the General tab on the Internet Options page, click Languages.

3. Select the language that you want to display:

Option	Description
Mozilla Firefox (ESR 31 and later)	Select the language that you want to display in the browser and click Move up until it's at the top of the list. Click Ok.
Google Chrome	Scroll to the Languages section, click the action customization next to the language that you want to display, and select Move to the top.
Internet Explorer (v11 and later)	Select the language that you want to display in the browser and click Move up until it's at the top of the list. Click Ok twice to return to the main browser page.

4. Refresh the browser tab where the IBM Spectrum Control GUI is displayed to apply the language setting.

Related tasks

- [Changing the language in a web browser doesn't change the entire IBM Spectrum Control GUI](#)

Changing the operating system language for Windows

Change the operating system language for Windows.

Procedure

To change the operating system language for Windows servers, for example from English to German (**de_DE**), complete these steps:

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start.

2. Click Control Panel .
3. Click Clock, Language, and Region > Region and Language.
Tip: The exact links that you click might vary between the different versions of Windows. For example, Region and Language might be separate links in more recent versions of Windows. In those cases, click Language.
4. On the Format tab, select a language and click OK.
Tip: The exact links that you click might vary between the different versions of Windows. For example, you might need to click the Administrative tab rather than the Format tab to change the system locale.
5. Restart all the IBM Spectrum® Control services or restart the Windows server.

Changing the operating system language for AIX

Change the operating system language for AIX®.

Procedure

To change the operating system language for AIX systems, for example from English to German (**de_DE**), complete these steps:

1. At the command prompt, run the following commands:

```
chlang -m de_DE de_DE
export LANG=de_DE
```

2. Restart all the IBM Spectrum® Control services or restart the AIX system.

Adding an installation license

If you installed IBM Spectrum® Control and want to add an IBM Spectrum Control license so that you can monitor storage systems with different license models, you can use the installation program or silent-mode installation.

About this task

To determine the type of license you have, find these files in the `<installation_dir>/swidtag` directory:

- ibm.com_IBM_Spectrum_Control-version.swidtag
Indicates that you have an IBM Spectrum Control license.
- ibm.com_IBM_Spectrum_Control_Standard_Select_Edition-version.swidtag
Indicates that you have an IBM Spectrum Control Select Edition license.

Depending on the license you purchased, one of the following license key files can be found in the **license\key** sub-directory of the electronic image:

node.lock.AE

IBM Spectrum Control license

node.lock.AS

IBM Spectrum Control Select Edition license

- [Adding an installation license using the installation program](#)
If you installed IBM Spectrum Control and want to add an IBM Spectrum Control license for additional function, you can do so by using the installation program.
- [Adding an installation license using silent mode](#)
If you installed IBM Spectrum Control and want to add an IBM Spectrum Control license to access more function, you can do so by using silent-mode installation.

Adding an installation license using the installation program

If you installed IBM Spectrum® Control and want to add an IBM Spectrum Control license for additional function, you can do so by using the installation program.

Procedure

To add the license to your system using the installation program, complete the following steps:

1. Log on to the IBM Spectrum Control system with the appropriate user privileges.
2. Start the IBM Spectrum Control installation program.
3. On the Choose Installation Location and Type page, select License upgrade and click Next.
4. On the Upgrade License page, select which license key file to use and click Next.
5. Verify the information on the Preinstallation Summary page and click Install.
6. On the Installation Completed page, click Done.
When the installer upgrades the license, it also stops and restarts the following IBM Spectrum Control servers: Web, Data, and Device.

Related reference

- [Planning for IBM Spectrum Control authentication and authorization](#)

Adding an installation license using silent mode

If you installed IBM Spectrum® Control and want to add an IBM Spectrum Control license to access more function, you can do so by using silent-mode installation.

Procedure

To add the license to your system that uses the silent-mode installation, complete the following steps:

1. Log on to your IBM Spectrum Control system with the appropriate user privileges.
2. Modify the following parameters in the silent_SingleServerTypical.properties, silent_SingleServerCustom.properties, or silent_MultipleServer.properties file:
 - CHOSEN_INSTALL_TYPE="License Upgrade"
 - varUseLicenseKeyOnImage=
 - varLicenseKeyFile=

For more information, see [Editing the response file](#).

3. Save the appropriate response file.
4. Start the IBM Spectrum Control installation program.
For Windows operating system, run the command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```

For AIX® or Linux® operating system, run the following command:

```
./setup.bin -l language -i silent -f absolute_path_to_response_file
```

When the installer upgrades the license, it also stops and restarts the IBM Spectrum Control Web, Data, and Device servers.

Related reference

- [Planning for IBM Spectrum Control authentication and authorization](#)

Installing Storage Resource agents

You can install Storage Resource agents by using the IBM Spectrum® Control user interface or a command.

About this task

You can select one of the following agent installation scenarios:

- To deploy a Storage Resource agent using the GUI, see [Deploying Storage Resource agents](#).
- To install the agent using a command, see [Installing Storage Resource agents by using a command](#).

Restriction:

- You cannot use Storage Resource agents as data sources for switches and fabrics. For Brocade switches, you must use CIM agents. For Cisco switches, you must use SNMP agents.
- Before you install or deploy a Storage Resource agent on a system, you must disable the firewall on that system.
- If you are using IBM® PowerHA® SystemMirror® for AIX®, a Storage Resource agent must be installed on each node of the cluster and all agents in a cluster must be configured to use the same listening port.
- To scan a cluster resource group, you must configure the cluster resource group to have at least one IP address that is accessible from the IBM Spectrum Control server.

To deploy a Storage Resource agent, you must log in as a user with the following authority:

On Windows operating systems

You must have Windows Administrator authority.

On UNIX or Linux® operating systems

You must be logged in as the root user.

- [Installing Storage Resource agents by using a command](#)

You typically install Storage Resource agents by using the IBM Spectrum Control GUI. However, if you must install Storage Resource agents locally, you can do so with limited support.

Installing Storage Resource agents by using a command

You typically install Storage Resource agents by using the IBM Spectrum® Control GUI. However, if you must install Storage Resource agents locally, you can do so with limited support.

About this task

During this installation method, a return code of zero means that a successful installation occurred and a nonzero return code means that an unsuccessful installation occurred. If you have an unsuccessful installation, you must review the log files to determine the problem.

The-force option can be used when you have a Storage Resource agent point to multiple servers. If an agent is installed on one server, and another server wants to install an agent in the same location, the second server can use the -force option to install the agent.

If you use this installation method, when you enter a directory to install the Storage Resource agent, do not add an ending slash mark (\). For example, do not specify C:\agent1\. This action causes the installation to fail.

If you run the agent as a non-daemon service (on-demand service), you must make sure that at least one protocol is valid for a successful connection from the server to the agent. See [Table 1](#) for the required parameters for each protocol.

Table 1. Parameters required for each protocol

Protocol	Description
SSH	Requires the user ID and password or user ID, certificate, and passphrase.
Windows (SMB)	Requires the user ID and password.
REXEC	Requires user ID and password.
RSH	Requires the user ID.

The Storage Resource Agent image contains the installation images for the Storage Resource agents in the following directory:

DVD_installation_image_location/data/sra/operating_system

See [Table 2](#) for the Storage Resource agent installation images.

Table 2. Storage Resource agent installation images

Operating system	Operating system name
AIX®	aix_power
Linux® x86	linux_ix86
Linux for Power Systems Servers	linux_power
Linux s390	linux_s390
Windows	windows

To install the Storage Resource agents locally, complete the following steps:

Procedure

1. Go to the installation image location:

```
cd DVD_installation_image_location
```

2. Run the following command:

```
bin/Agent -install -serverPort server_secure_port -force
-serverIP server_IP_address -installLoc Agent_install_location -debug MAX
```

```

>> -agentPort agent_port — -commtype daemon >>
>> 1 >>
>> -userID user_ID — -password password — -certFile certificate_file >>
>> 2 >>
>> -passphrase passphrase >>
>> 3 >>

```

Notes:

- ¹ Parameters when the agent is run as a daemon service.
- ² Parameters when the agent is run as a non-daemon service. See [Table 1](#) to determine which parameter is required for each protocol.
- ³ Parameters when the agent is run as a non-daemon service. See [Table 1](#) to determine which parameter is required for each protocol.

The parameters are as follows.

```

-install
    Installs the Storage Resource agent.

-force
    Forces the agent to be installed. The two different situations in which this parameter must be specified are as follows.
    • If an earlier installation failed and there is residue on the system that causes further installations to fail. You must verify that all the parameters that are provided are valid, such as the installation location, port .
    • If the agent is already installed from one server and you now must install the agent pointing to another server.

-serverPort server_secure_port
    The default port for the Data server is 9549.

-serverIP server_IP_address
    The IP address of the server. If the server can be reached through multiple IP addresses, then multiple IP addresses can be specified with IP addresses separated with a comma.

-installLoc "Agent_install_location"
    Location where the agent is installed. Enclose the directory name in quotation marks, for example, "C:\Program Files\IBM\TPC_SRA".

-debug MAX
    Optional parameter for debugging purposes.

-agentPort agent_port
    If the agent is run as a daemon service, the agent port must be specified. The default agent port is 9510.

-commtype daemon
    If the agent is run as a daemon service, then this parameter must be specified.

-userID user_ID
    For non-daemon service. The user ID defined on the agent system. This user ID is used by the server to connect to the agent system.

-password password
    For non-daemon service. Password for the user ID.

-certFile "certificate_file"
    For non-daemon service. The certificate that is used for SSH communication between the server and agent. This certificate must be stored on the server system. Enclose the directory name in quotation marks, for example, "c:\keys\id_sra".
    For information about SSH protocol and SSL protocol certificates, see Creating a certificate for SSH protocol and Replacing default SSL certificates for the Data server and Storage Resource agents with custom SSL certificates.

-passphrase passphrase
    For non-daemon service. The passphrase that is defined for the certificate that is used in SSH communication.

```

What to do next

After you install a Storage Resource agent by using a command, you must schedule a probe to collect data about the associated resource. For information about how to schedule a probe, see [Creating probes](#).

Installing IBM Cognos Analytics

You can install the optional IBM® Cognos® Analytics 11.2.0 or later reporting tool on one computer, on multiple servers for a distributed installation, or you can expand an existing single computer installation to another server to improve performance. Cognos Analytics provides an installation program that guides you through the installation process.

Once you have installed Cognos Analytics you can view predefined reports and create custom reports about IBM Spectrum® Control. You can access reports from IBM Spectrum Control, and work with the reports in the Cognos Analytics reporting tool.

- [Installing Cognos Analytics in a single Windows environment](#)
You can install Cognos Analytics in a Microsoft Windows environment by using the Cognos Analytics installation program.
- [Installing Cognos Analytics in a single Linux environment](#)
You can install Cognos Analytics in a Linux® environment by using the Cognos Analytics installation program. To run the graphical-mode installation program, your computer must support a Java™-based graphical user interface.
- [Installing Cognos Analytics in a single AIX environment](#)
You can install Cognos Analytics in an AIX® environment by using the Cognos Analytics installation program. To run the graphical-mode installation program, your computer must support a Java-based graphical user interface.

Installing Cognos Analytics in a single Windows environment

You can install Cognos® Analytics in a Microsoft Windows environment by using the Cognos Analytics installation program.

Before you begin

The IBM® Cognos Analytics 11.2.0 or later installation consists of the following components:

- Installer executable file
- Repository .zip file

Important: When you run the installer, you must point to the server repository.zip file.

For Microsoft Windows operating system installations, ensure that you have administrator privileges for the Windows computer you are doing the installation on. Also, ensure that your computer has a TEMP system variable that points to the directory where you want to store temporary files. During installation, files from the disk are temporarily copied to this directory.

For more information, see [Verify system requirements](#).

For more information, see [Software Product Compatibility Reports](#).

Procedure

1. Log in to your Windows computer as a user with administrative privileges.
2. Start the Cognos Analytics installer executable file.
3. Select the language to use for your installation and click Next.
4. On the Repository selection page, click Choose.
5. On the Please Choose a File page, select the server repository zip file and click Open.
6. Click Next.
7. On the Product Install page, select IBM Cognos Analytics and click Next.
8. On the License Agreement page, accept the license agreement and click Next.
9. On the Location page, provide your installation location and name of the Shortcut folder and click Next.
If you want the IBM Cognos Analytics installation to automatically create the installation location for you, click Yes.
10. On the Installation type page, select Custom and click Next.
The Custom option offers full flexibility to select the IBM Cognos Analytics components that you want to install. You might want to customize or integrate Cognos Analytics with third-party software? If so, this option is one you would select.
11. On the Choose components page, accept the default selections, and click Next.
12. On the Pre-Installation Summary page, review your information for accuracy and click Install.
13. When the installation is complete, click Done.

Related information

- [Silent installation, uninstallation, and configuration](#)
- [Uninstall IBM Cognos Analytics on Microsoft Windows operating systems](#)

Installing Cognos Analytics in a single Linux environment

You can install Cognos® Analytics in a Linux® environment by using the Cognos Analytics installation program. To run the graphical-mode installation program, your computer must support a Java™-based graphical user interface.

Before you begin

The IBM® Cognos Analytics 11.2.0 or later installation consists of the following components:

- Installer executable file
- Repository .zip file

When you run the installer, you must point to the server repository .zip file.

For Linux operating system installations, ensure that you are a root user for the computer you are doing the installation on.

Verify that the Cognos Analytics installer executable file has 755 permissions and that the Cognos Analytics repository .zip file has 644 permissions.

You *do not* need to set the JAVA_HOME environment variable before you start the IBM Cognos Analytics 11.2.0 or later installation program.

Verify that you have installed the following prerequisite packages on a Linux computer for IBM Cognos Analytics 11.2.0 or later.

The prerequisite packages for Red Hat® Enterprise Linux (RHEL) 6 are:

```
glibc-2.12-1.166.el6 or later (both i686 and x86_64 packages)
libstdc++-4.4.7-16.el6 or later (both i686 and x86_64 packages)
nspr-4.9.2-1.el6 or later (both i686 and x86_64 packages)
nss-3.14.0-12.el6 or later (both i686 and x86_64 packages)
openmotif-2.3.3-5.el6 or later (both i686 and x86_64 packages)
```

The prerequisite packages for Red Hat Enterprise Linux (RHEL) 7 are:

```
glibc-2.17-55.el7 or later (both i686 and x86_64 packages)
libstdc++-4.8.2-16.el7 or later (both i686 and x86_64 packages)
nspr-4.10.2-4.el7 or later (both i686 and x86_64 packages)
nss-3.15.4-6.el7 or later (both i686 and x86_64 packages)
motif-2.3.4-7.el7 or later (both i686 and x86_64 packages)
```

The prerequisite packages for Red Hat Enterprise Linux (RHEL) 8 are:

```
unzip-6.0-41.el8.x86_64 or later
libstdc++-8.3.1-4.5.el8 or later (both i686 and x86_64 packages)
```

```
glibc-2.28-72.el8 or later (both i686 and x86_64 packages)
libXaw-1.0.13-10.el8.x86_64 or later
libXtst-1.2.3-7.el8.x86_64 or later
libXi-1.7.9-7.el8.x86_64 or later
libXft-2.3.2-10.el8 or later (both i686 and x86_64 packages)
libX11-1.6.7-1.el8 or later (both i686 and x86_64 packages)
motif-2.3.4-16.el8 or later (both i686 and x86_64 packages)
nss-3.44.0-9.el8_1 or later (both i686 and x86_64 packages)
nspr-4.21.0-2.el8_0 or later (both i686 and x86_64 packages)
libnsl-2.28-72.el8 or later (both i686 and x86_64 packages)
```

For more information, see [Verify system requirements](#).

For more information, see [Software Product Compatibility Reports](#).

You need to set the correct ulimit values in the /etc/security/limits.conf file for IBM Cognos Analytics 11.2.0 or later.

Here is an example of the entries you might need to add to the /etc/security/limits.conf file:

```
*      soft    nproc    unlimited
*      hard    nproc    unlimited
*      soft    nofile   8192
*      hard    nofile   8192
*      soft    stack    unlimited
*      hard    stack    unlimited
```

After you set and save your ulimit values, restart your computer and verify the correct ulimit values.

For more information, see [Set the ulimit values on UNIX and Linux operating systems](#).

Procedure

1. Log in to your Linux computer as a root user.
2. Start the Cognos Analytics installer executable file.
3. Select the language to use for your installation and click Next.
4. On the Repository selection page, click Choose.
5. On the Please Choose a File page, select the server repository zip file and click Open.
6. Click Next.
7. On the Product Install page, select IBM Cognos Analytics and click Next.
8. On the License Agreement page, accept the license agreement and click Next.
9. On the Location page, provide your installation location and click Next.
If you want the IBM Cognos Analytics installation to automatically create the installation location for you, click Yes.
10. On the Choose components page, accept the default selections, and click Next.
11. On the Pre-Installation Summary page, review your information for accuracy and click Install.
12. When the installation is complete, click Done.

Related information

- [Silent installation, uninstallation, and configuration](#)
- [Uninstall IBM Cognos Analytics on UNIX or Linux operating systems](#)

Installing Cognos Analytics in a single AIX environment

You can install Cognos® Analytics in an AIX® environment by using the Cognos Analytics installation program. To run the graphical-mode installation program, your computer must support a Java™-based graphical user interface.

Before you begin

The IBM® Cognos Analytics 11.2.0 or later installation consists of the following components:

- Installer executable file
- Repository .zip file

When you run the installer, you must point to the server repository .zip file.

For AIX operating system installations, ensure that you are a root user for the computer you are installing on.

Verify that the Cognos Analytics installer executable file has 755 permissions and that the Cognos Analytics repository .zip file has 644 permissions.

You *do not* need to set the JAVA_HOME environment variable before you start the Cognos Analytics 11.2.0 or later installation program.

Verify that you have the following prerequisite packages installed on a AIX computer for Cognos Analytics 11.2.0 or later.

The prerequisites for AIX 7.1 or later are:

- IBM XL C/C++ for AIX, V12.1 Runtime Environment. If using AIX XL C++ Runtime it is recommended to use 13.1.2.160112 or later.
- bos.pmap.lib

For more information, see [Verify system requirements](#).

For more information, see [Software Product Compatibility Reports](#)

You need to set the correct ulimit values in the /etc/security/limits file for IBM Cognos Analytics 11.0.9 or later.

Here is an example of the entries you might need to add to the `/etc/security/limits` file:

```
default:
fsize = -1
core = -1
cpu = -1
rss = -1
stack = 16384
nofiles = 8192
```

After you set and save your ulimit values, restart your computer and verify the correct ulimit values.

For more information, see [Set the ulimit values on UNIX and Linux® operating systems](#).

Procedure

1. Log in to your AIX computer as a root user.
2. Start the Cognos Analytics installer executable file.
3. Select the language to use for your installation and click Next.
4. On the Repository selection page, click Choose.
5. On the Please Choose a File page, select the server repository zip file and click Open.
6. Click Next.
7. On the Product Install page, select IBM Cognos Analytics and click Next.
8. On the License Agreement page, accept the license agreement and click Next.
9. On the Location page, provide your installation location and click Next.
If you want the IBM Cognos Analytics installation to automatically create the installation location for you, click Yes.
10. On the Choose components page, accept the default selections, and click Next.
11. On the Pre-Installation Summary page, review your information for accuracy and click Install.
12. When the installation is complete, click Done.

Related information

- [Silent installation, uninstallation, and configuration](#)
- [Uninstall IBM Cognos Analytics on UNIX or Linux operating systems](#)

Configuring IBM Cognos Analytics

After you install the optional IBM® Cognos® Analytics, to view predefined reports and create custom reports about IBM Spectrum® Control, you must copy Db2® files, create a content store database, save your configuration settings, and start the services.

Important: The configuration procedures assume that you are going to use Db2 for the Cognos Analytics content store. *It is not required* that the Cognos Analytics content store use Db2. Cognos Analytics supports other database products for the content store.

The IBM Spectrum Control database repository uses Db2. If you do not use Db2 for the content store, you still need to copy the Db2 JDBC driver files into the Cognos Analytics installation location and link the 32-Bit Db2 library file (Linux® and AIX® only) in order for Cognos Analytics to connect to the IBM Spectrum Control database repository.

A general overview to complete the configuration tasks in order to view predefined reports and create custom reports about IBM Spectrum Control:

- Copying Db2 files and linking Db2 libraries
- Creating the content store
- Creating a data source
- Importing pre-defined reports package
- [Copying JDBC files for Db2 on Windows](#)
You must copy JDBC driver files from your Db2 installation directory to your Cognos Analytics installation location.
- [Copying Db2 files and linking the Db2 library on Linux and AIX](#)
You must copy JDBC driver files from your Db2 installation directory to your Cognos Analytics installation location and then link the Db2 library.
- [Creating a content store and starting Cognos Analytics on Windows](#)
The information that you enter in IBM Cognos Configuration for Cognos Analytics is used to generate a script that you use to create the content store database.
- [Creating a content store and starting Cognos Analytics on Linux or AIX](#)
The information that you enter in IBM Cognos Configuration for Cognos Analytics is used to generate a script that you use to create the content store database.
- [Creating Cognos Analytics data source for the IBM Spectrum Control database](#)
You can create an IBM Cognos Analytics data source for your IBM Spectrum Control database.
- [Importing IBM Spectrum Control pre-defined reports package](#)
Complete this procedure to import the IBM Spectrum Control predefined reports package into the IBM Cognos Analytics reporting tool.
- [Configuring access to the Cognos Analytics reporting tool](#)
Specify which user IDs and groups of user IDs can use reports, report tools, and report administration functions in the Cognos Analytics tool.

Related information

- [Guidelines for creating the content store](#)
- [Set up database connectivity for the content store database](#)
- [Set Database Connection Properties for the content store](#)

Copying JDBC files for Db2 on Windows

You must copy JDBC driver files from your Db2® installation directory to your Cognos® Analytics installation location.

Procedure

Copy the following Db2 JDBC driver files from the *Db2_installation_dir\java* directory to the *Cognos_installation_dir\drivers* directory:

- db2jcc.jar - This JDBC driver file does not exist for Db2 11.5 or later
- db2jcc_license_cu.jar
- db2jcc4.jar

For more information, see [Configuration actions that are critical to the success of your installation](#).

Copying Db2 files and linking the Db2 library on Linux and AIX

You must copy JDBC driver files from your Db2® installation directory to your Cognos® Analytics installation location and then link the Db2 library.

Procedure

1. Copy the following Db2 JDBC driver files from the *Db2_installation_dir/java/* directory to the *Cognos_installation_dir/drivers/* directory:
 - db2jcc.jar - This JDBC driver file does not exist for Db2 11.5 or later
 - db2jcc_license_cu.jar
 - db2jcc4.jar

For more information, see [Configuration actions that are critical to the success of your installation](#).

2. Enter the following command as the root user:

```
cd Cognos_installation_dir/drivers/  
chmod +x db2jcc*.jar
```

3. Enter the following command as the root user:
On the Linux® operating system:

```
ln -s Db2_instance_owner_home_directory/sqlib/lib32/libdb2.so Cognos_installation_dir/bin/
```

On the AIX® operating system:

```
ln -s Db2_instance_owner_home_directory/sqlib/lib32/libdb2.a Cognos_installation_dir/bin/
```

In the IBM Spectrum® Control Predefined Reports use Compatibility Query Mode, Cognos Analytics needs to use the 32-Bit Db2 library in order to connect to the IBM Spectrum Control database as a data source.

Creating a content store and starting Cognos Analytics on Windows

The information that you enter in IBM® Cognos® Configuration for Cognos Analytics is used to generate a script that you use to create the content store database.

Procedure

1. Create the content store for Cognos Analytics.
 - a. From the Windows Start menu, select IBM Cognos Analytics > IBM Cognos Configuration.
 - b. In the Explorer window, expand Data Access > Content Manager and click Content Store.
 - c. In the Type property, enter DB2 database.
 - d. In the Database server and port number property, enter localhost or the fully qualified domain name of the machine where Db2® is installed.
For example, *localhost:25000* or *FQDN_Of_Your_Machine:25000* where 25000 is the default port number that is used by Db2. If you are using a different port number (for example, 50000), ensure that you use that value. To determine the port that is used by Db2, examine the varDBPort parameter in the *installation_dir\config\InstallVariable.properties* file.
 - e. In the User ID and password property, enter the user name of the owner of the Db2 instance where the content store database will be located and the password that is associated with that user name.
Typically, this user name is *db2admin*. If the user name is a Windows domain user name, enter the user name *without* a Windows domain name prefix or suffix.
 - f. In the Database name property, verify that the default value is cm.
 - g. From the File menu, click Save.
 - h. In the informational window, click Close.
 - i. In the Explorer window, expand Data Access > Content Manager and click Content Store.
 - j. Right-click Content Store and click Generate DDL.
 - k. In the informational window, click Close.
 - l. From the Windows Start menu, select IBM Db2 Db2COPY1 (Default) > Db2 Command Window - Administrator. Then, enter the following command:

```
db2 -tvf "Cognos_installation_dir\configuration\schemas\content\db2\createDb.sql"
```

On Windows, add quotes around the directory path if it contains spaces. For example:

```
db2 -tvf "C:\Program Files\IBM\Cognos\analytics\configuration\schemas\content\db2\createDb.sql"
```

This process might take a few minutes.

- m. From the Actions menu, click Test.
You are testing the connection between content manager and the content store database.

- n. In the informational window, click Close.
For more information, see [Generating a script file to create a database for a Db2 content store](#).
2. Start Cognos Analytics (content manager and application services).
It might take Cognos Analytics a few minutes to start.
 - a. In IBM Cognos Configuration, in the Explorer window, click Local Configuration.
 - b. From the Actions menu, click Test.
Tip: If you are not configuring Cognos Analytics to send notifications, you can ignore the testing the mail server connection failure message. If there is a failure related to NIST SP800-131A conformance, click Details in the dialog to see more information. The integration between Cognos Analytics and IBM Spectrum® Control is not affected by this setting.
 - c. From the Actions menu, click Start.
3. Exit IBM Cognos Configuration.

Creating a content store and starting Cognos Analytics on Linux or AIX

The information that you enter in IBM® Cognos Configuration for Cognos Analytics is used to generate a script that you use to create the content store database.

Procedure

1. Create the content store for Cognos Analytics.
 - a. Enter the following commands to start IBM Cognos Configuration:


```
cd <Cognos_installation_dir>/bin64/
. /home/db2inst1/sqllib/db2profile
./cogconfig.sh
```

Notice the space between the . and the / in the second command.

Tip:
It's recommended to source-in the db2profile (. /home/db2inst1/sqllib/db2profile) anytime you are going to start IBM Cognos Configuration.
 - b. In the Explorer window, expand Data Access, Content Manager and click Content Store.
 - c. In the Database server and port number property, enter localhost or the fully qualified domain name of the computer where Db2® is installed.
For example, localhost:25000 or FQDN_Of_Your_Machine:25000 where 25000 is the default port number that is used by Db2. If you are using a different port number (for example, 50000), ensure that you use that value. To determine the port that is used by Db2, examine the varDBPort parameter in the installation_dir/config/InstallVariable.properties file.
 - d. In the User ID and password property, enter the user name of the owner of the Db2 instance where the content store database is located and the password that is associated with that user name.
 - e. In the Database name property, verify that the default value is cm.
 - f. From the File menu, click Save.
 - g. In the Explorer window, go to Data Access, Content Manager and click Content Store.
 - h. Right-click Content Store and click Generate DDL.
 - i. In the informational window, click Close.
 - j. While you are logged in to the Linux® or AIX® operating system as the root user, switch to the user that is the Db2 instance owner (for example, db2inst1).
 - k. Enter the following Db2 command:

```
db2 -tvf Cognos_installation_dir/configuration/schemas/content/db2/createDb.sql
```

If the results of this command contain the following similar messages, you can safely ignore those messages because the Db2 instance owner already has the necessary privileges:

```
GRANT CREATETAB,BINDADD,CONNECT,IMPLICIT_SCHEMA ON DATABASE TO USER db2inst1
DB21034E The command was processed as an SQL statement because it was not a
valid Command Line Processor command. During SQL processing it returned:
SQL0554N An authorization ID cannot grant a privilege or authority to itself.
SQLSTATE=42502
```

```
GRANT CREATEIN,DROPIN,ALTERIN ON SCHEMA DB2COGNOS TO USER db2inst1
DB21034E The command was processed as an SQL statement because it was not a
valid Command Line Processor command. During SQL processing it returned:
SQL0554N An authorization ID cannot grant a privilege or authority to itself.
SQLSTATE=42502
```

```
GRANT USE OF TABLESPACE TSN_USR_cm TO USER db2inst1
DB21034E The command was processed as an SQL statement because it was not a
valid Command Line Processor command. During SQL processing it returned:
SQL0554N An authorization ID cannot grant a privilege or authority to itself.
SQLSTATE=42502
```

```
GRANT USE OF TABLESPACE TSN_REG_cm TO USER db2inst1
DB21034E The command was processed as an SQL statement because it was not a
valid Command Line Processor command. During SQL processing it returned:
SQL0554N An authorization ID cannot grant a privilege or authority to itself.
SQLSTATE=42502
```

- l. Go back to the root user.
- m. From the Actions menu, click Test.
You are testing the connection between Content Manager and the content store database.
- n. In the informational window, click Close.
For more information, see [Generating a script file to create a database for a Db2 content store](#).
2. Start Cognos Analytics (Content Manager and application services).
 - a. In IBM Cognos Configuration, in the Explorer window, click Local Configuration.

- b. From the Actions menu, click Test.
Tip: If you are not configuring Cognos Analytics to send notifications, you can ignore the testing the mail server connection failure message. If there is a failure related to NIST SP800-131A conformance, click Details in the dialog to see more information. The integration between Cognos Analytics and IBM Spectrum® Control is not affected by this setting.
 - c. From the Actions menu, click Start.
3. When Cognos Analytics starts, exit IBM Cognos Configuration.

Creating Cognos Analytics data source for the IBM Spectrum Control database

You can create an IBM® Cognos® Analytics data source for your IBM Spectrum® Control database.

Before you begin

The following table shows the Db2® actions that are required before you create an IBM Cognos Analytics data source for the IBM Spectrum Control database:

Table 1. Required Db2 actions for IBM Spectrum Control/Cognos Analytics integration scenarios

IBM Spectrum Control - Cognos Analytics integration scenario	Required Db2 actions
IBM Spectrum Control database and Cognos Analytics on the same Windows or Linux® computer	There are no Db2 actions required
IBM Spectrum Control database and Cognos Analytics on the same AIX® computer	<p>Configure the connectivity with the IBM Spectrum Control database to use the loopback mode, or Cognos Analytics will fail to connect to the IBM Spectrum Control database source when running reports multiple times.</p> <p>Complete the following steps on the AIX computer:</p> <ol style="list-style-type: none"> 1. Log on to the computer as a user with root privileges 2. Switch to the user that is the Db2 instance owner (for example, <i>db2inst1</i>). 3. Run this command: <pre>db2 catalog tcpip node loopbk remote 127.0.0.1 server 25000</pre> <p>If your Db2 instance uses a different port other than 25000, provide that port value in the command. For example, 50000.</p> 4. Run this command: <pre>db2 catalog database TPCDB as TPCDB_LP at node loopbk</pre> <p>If your IBM Spectrum Control database has a name other than TPCDB, provide that database name in the command.</p> <p>If you have upgraded from an earlier version of IBM Spectrum Control, you might find that you already have the loopback connection configured and the above mentioned commands return errors. Confirm the configuration using the 'db2 list database directory' command and look for TPCDB_LP.</p>
IBM Spectrum Control database on a different computer than where you installed Cognos Analytics, all platforms	<p>Catalog the node and database for the remote IBM Spectrum Control database.</p> <p>Complete the following steps on the computer where you installed Cognos Analytics:</p> <ol style="list-style-type: none"> 1. On Windows operating system: Log in as a user with administrative privileges. <ol style="list-style-type: none"> a. On Linux or AIX operating system: Log in as a user with root privileges. 2. On Windows operating system, from the Start Menu, select IBM DB2 DB2COPY1 (Default) > DB2 Command Window - Administrator. <ol style="list-style-type: none"> a. On Linux or AIX operating system: Switch to the user that is the Db2 instance owner (for example, <i>db2inst1</i>). 3. Enter the following command to catalog the remote node: <pre>db2 catalog tcpip node node_name remote server_fqdn server port</pre> <p>Where <i>node_name</i> is a new value you provide to name the node. For example, remoteSC. The <i>server_fqdn</i> value is the fully qualified domain name of the machine where the IBM Spectrum Control database resides. For example, myserver.mycompany.com. The <i>port</i> value is the port number that is used by Db2. The default value is 25000. Port 50000 can also be used.</p> <p>Example:</p> <pre>db2 catalog tcpip node remoteSC remote myserver.mycompany.com server 25000</pre> 4. Enter the following command to catalog the remote database: <pre>db2 catalog database database_name as alias at node node_name</pre> <p>Where <i>database_name</i> is the name of the IBM Spectrum Control database. Typically, this name is TPCDB. The <i>alias</i> value is a new value that you provide to reference the cataloged database. For example, remDB. The <i>node_name</i> value is the value that you set in Step 3.</p> <p>Example:</p> <pre>db2 catalog database TPCDB as remDB at node remoteSC</pre>

Procedure

To create an Cognos Analytics data source for the IBM Spectrum Control database, complete the following steps:

1. Open a web browser and enter the following URL:

`http://FQDN_Of_Your_Machine:9300/bi`

Where **http** is the protocol, and *FQDN_Of_Your_Machine* is your fully qualified domain name, (for example, *myserver.mycompany.com*).

2. Select Manage Administration console.
3. In the IBM Cognos Administration GUI, select the Configuration tab.
4. In the Configuration tab, select Data Source Connections.
5. On the toolbar, click the New Data Source icon. It is the first icon in the list.
6. On the Specify a name and description page, in the Name property, enter **TPCDB**.
You must enter the TPCDB value or the IBM Spectrum Control predefined reports will not generate successfully.
7. On the Specify a name and description page, click Next.
8. On the Specify the connection page, select IBM DB2 from the Type field.
9. Verify that Use the default object gateway is selected.
10. Unmark Configure JDBC connection and click Next.
11. On the Specify the IBM DB2 connection string page, in the DB2 database name property, complete one of the following steps:
 - a. If the IBM Spectrum Control database is on the same Windows or Linux computer as where you installed Cognos Analytics (first row scenario from Table 1), enter the name of the IBM Spectrum Control database. Typically, this name is **TPCDB**.
 - b. If the IBM Spectrum Control database is on the same AIX computer as where you installed Cognos Analytics (second row scenario from Table 1), enter the alias value that you created when you cataloged the loopback database. Typically, this value is **TPCDB_LP**.
 - c. If the IBM Spectrum Control database is on a different computer than where you installed Cognos Analytics (third row scenario from Table 1), enter the *alias* value that you created when you cataloged the remote database. For example, **remDB**.
12. On the Specify the IBM DB2 connection string page, in the Signon section, do the following:
 - a. Verify that Signons is selected.
 - b. Select Password.
 - c. In the User ID property, enter the user name of the owner of the Db2 instance where the IBM Spectrum Control database is located. On a Window operating system, this user name is typically *db2admin*.
On Linux or AIX operating system, this user name is typically *db2inst1*.
 - d. In the Password and Confirm password properties, enter the password that is associated with the User ID property from the previous substep.
13. On the Specify the IBM DB2 connection string page, in the Testing section, click Test the connection.
14. On the Test the connection page, click Test.
15. When the test is successful, on the View the results page, click Close.
16. On the Test the connection page, click Close.
17. On the Specify the IBM DB2 connection string page, click Next.
18. On the Specify the commands page, click Finish.
19. Close your web browser.

- [Using Cognos Analytics with multiple remote IBM Spectrum Control databases](#)

You can configure multiple connections within a single Cognos Analytics data source to use Cognos Analytics with multiple remote IBM Spectrum Control databases.

Related information

- [Creating a data source connection](#)

Using Cognos Analytics with multiple remote IBM Spectrum Control databases

You can configure multiple connections within a single Cognos® Analytics data source to use Cognos Analytics with multiple remote IBM Spectrum® Control databases.

Before you begin

On the machine where you have Cognos Analytics installed, you must catalog the node and the database for each remote IBM Spectrum Control database that you want to query when generating Cognos Analytics reports.

Follow the steps in the *Before you begin* section of [Creating Cognos Analytics data source for the IBM Spectrum Control database](#) to catalog the node and the database for each remote IBM Spectrum Control database. Then, follow the steps in the *Procedure* section of the same topic to create the single data source for the IBM Spectrum Control databases and to create the first connection to a specific IBM Spectrum Control database within that data source.

At this point, you will have one data source named **TPCDB** defined in Cognos Analytics and within that one data source, you will have one connection defined, which is also named **TPCDB**.

About this task

It's assumed that Cognos Analytics is installed on its own machine and that you have multiple IBM Spectrum Control systems installed on separate machines.

Procedure

To rename the **TPCDB** connection and to create more connections to other remote IBM Spectrum Control databases within the single data source, complete the following steps:

1. Open a web browser and enter the following URL:

`http://FQDN_Of_Your_Machine:9300/bi`

Where `http` is the protocol, and `FQDN_Of_Your_Machine` is your fully qualified domain name, (for example, `myserver.mycompany.com`) of the machine where you installed Cognos Analytics.

2. Select Manage Administration console.
3. In the IBM Cognos Administration GUI, select the Configuration tab.
4. In the Configuration tab, select Data Source Connections.
5. On the Data Source Connections page, click the TPCDB data source. This action will drill down one level and you will see the TPCDB connection.
6. In Actions, click the Set properties - TPCDB icon; it's the first icon in the list.
7. On the Set properties - TPCDB page, change the Name value from TPCDB to a value that reflects the particular IBM Spectrum Control database that this connection uses (for example, `scserv1`).
8. Click OK.
You are returned to the Data Source Connections page and your first data source has been renamed.
9. On the Data Source Connections page, in the toolbar, click the New Connection icon. It's the first icon in the list.
10. On the Specify a name and description - New Connection page, in the Name property, enter a value that reflects the particular IBM Spectrum Control database that this connection will use (for example, `myserv2`).
11. On the Specify a name and description - New Connection page, click Next.
12. On the Specify a name and description - New Connection page, select IBM Db2 from the Type field.
13. Verify that Use the default object gateway is selected.
14. Unmark Configure JDBC connection and click Next.
15. On the Specify the IBM Db2 connection string - New Connection page, in the DB2 database name property, enter the *alias* value that you created when you cataloged a particular remote IBM Spectrum Control database.
16. On the Specify the IBM Db2 connection string - New Connection page, in the Signon section, do the following:
 - a. Verify that Signons is selected.
 - b. Select Password.
 - c. In the User ID property, enter the user name of the owner of the Db2® instance where the remote IBM Spectrum Control database is located. Note: On a Windows operating system, this user name is typically `db2admin`. On Linux® or AIX® operating systems, this user name is typically `db2inst1`.
 - d. In the Password and Confirm password properties, enter the password that is associated with the User ID value from the previous substep.
17. On the Specify the IBM Db2 connection string - New Connection page, in the Testing section, click Test the connection.
18. Click Test.
19. When the test is successful, on the View the results page, click Close.
20. On the Test the connection - New Connection page, click Close.
21. On the Specify the IBM Db2 connection string - New Connection page, click Next.
22. On the Specify the commands - New Connection page, click Finish.

Results

Your new connection is added to the list of connections defined for the single **TPCDB** data source. You can repeat Steps 10 - 23 to create a new connection for each remote IBM Spectrum Control database that you want to use when you generate reports. When you have multiple connections defined and you generate a report, you are prompted to select a specific connection to use for that report.

Importing IBM Spectrum Control pre-defined reports package

Complete this procedure to import the IBM Spectrum® Control predefined reports package into the IBM® Cognos® Analytics reporting tool.

Procedure

1. Copy the IBM Spectrum Control predefined reports package `TPC_deployment_pkg.zip` file from the `installation_dir\data\gui\` directory, on the machine where you installed IBM Spectrum Control, to the `Cognos_Install_Dir\deployment\` directory on the machine where you installed Cognos Analytics.

2. Import the package into Cognos Analytics.
 - a. Open a web browser and enter the following URL:

`http://machine_FQDN:9300/bi`

Where `machine_FQDN` is the fully qualified domain name of the machine where you installed Cognos Analytics. For example, `mycognos.mycompany.com`.
Warning: At this point Cognos Analytics is configured with anonymous access so that anyone is able to access it.

- b. Select Manage Administration console.
- c. In the IBM Cognos Administration GUI, select the Configuration tab.
- d. In the Configuration tab, select Content Administration.
- e. On the toolbar, click the New Import icon.
It is the fourth icon in the list.
- f. On the Select a deployment archive page, select `TPC_deployment_pkg` and click Next.
- g. On the Specify a name and description page, enter a name for your import and click Next.
- h. On the Select the public folders, directory, and library content page, select the following entries:
 - IBM Spectrum Control Packages
 - IBM Spectrum Control Predefined Reports
 - IBM Spectrum Control Report Layouts
- i. Click Next.
- j. On the Specify the general options page, click Next.
- k. On the Review the summary page, click Next.
- l. On the Select an action page, select Save and run once and click Finish.
- m. On the Run with Options page, select Upgrade all report specifications to the latest version and Assign new IDs during import.
For more information, see [Run an import](#).
- n. Click Run.
- o. On the IBM Cognos software page, select View the details of this import after closing this dialog and click OK.

- p. On the View run history details page, click Refresh.
Do *not* log off from the IBM Cognos Administration GUI until the import job is complete.
Important:
The import job might take 15 minutes or more to resolve. Click Refresh periodically to check the status.
- q. After the import job is complete, on the View an import deployment record page, click Close.
- r. Exit the IBM Cognos Administration GUI and return to the Cognos Analytics GUI.
3. Verify that the imported content is displayed in Cognos Analytics.
- a. Open a web browser and enter the following URL:

`http://machine_FQDN:9300/bi`

Where *machine_FQDN* is the fully qualified domain name of the machine where you installed Cognos Analytics. For example, *mycognos.mycompany.com*.

b. On the IBM Cognos Analytics Welcome page, click Team content.

c. Verify that the following content is listed:

- IBM Spectrum Control Packages
- IBM Spectrum Control Predefined Reports
- IBM Spectrum Control Report Layouts

Configuring access to the Cognos Analytics reporting tool

Specify which user IDs and groups of user IDs can use reports, report tools, and report administration functions in the Cognos® Analytics tool.

About this task

You can configure access in the following ways in the Cognos Analytics reporting tool:

- Create user IDs and groups of user IDs so that you can specify access for those users and groups.
- Assign administrative privileges to administrator IDs and remove administrative privileges from other user IDs.
- Specify which user IDs or groups can access individual reports or report folders.
- Specify which user IDs or groups can access individual Cognos Analytics reporting tool functions.

Information about how to configure security for Cognos Analytics is available in the IBM® Docs for Cognos Analytics. Use the following links to view this information:

Table 1. Information about configuring security for Cognos Analytics

Task	Link to information
Learn about the security model, authorization, and the default Cognos namespace in Cognos Analytics.	Security model
Learn about the initial security settings in Cognos Analytics.	Initial security
Learn about the initial access permissions in Cognos Analytics.	Initial access permissions
Learn about users, groups, and roles in Cognos Analytics.	Users, Groups, and Roles
Learn about configuring Cognos Analytics with LDAP authentication.	Configuring authentication providers
Learn about configuring access permissions in Cognos Analytics.	Access Permissions and Credentials
	Set access permissions for an entry

Reinstalling the software if a failure occurs

If an installation failure occurs, you do not have to uninstall components that were successfully installed. IBM Spectrum® Control provides an option to partially, or fully, roll back the installation.

A full rollback uninstalls all the components (even if they were installed successfully), and a partial rollback uninstalls only the components that were not successfully installed. The partial rollback option helps you resolve the problems that are causing the installation failure. After you resolve the problem, you can install the remaining components. For example, if the correct libraries were not installed for the Linux® operating system, and the installation stops. You can install the correct libraries and resume the installation.

If you install IBM Spectrum Control on the Windows operating system and select Full Rollback, you must restart the server after the rollback is complete.

If you originally selected a partial rollback, but decided to completely uninstall, you must run the uninstallation program before you reinstall IBM Spectrum Control.

The following table shows the results from a partial rollback. The components are listed in the order in which they were installed.

Table 1. Results from a partial rollback

If an installation failure occurs when you install this component...	Partial rollback results are...
IBM Spectrum Control common files and Java™ Runtime Environment	There is no partial rollback. The installation program completes a full rollback.
Database repository	The database repository is uninstalled.
Data server	The Data server is uninstalled.
WebSphere® Application Server Liberty profile	WebSphere Application Server Liberty profile is uninstalled.
Alert server	The Alert server is uninstalled.
Device server	The Device server is uninstalled.
Export server	The Export server is uninstalled.
IBM Spectrum Control GUI	The IBM Spectrum Control GUI is uninstalled.
CLI	The command-line interface is uninstalled.
Storage Resource agent	The Storage Resource agent is uninstalled.

Taking the first steps after installation

After IBM Spectrum® Control is installed, configure it to monitor the resources in your environment.

Access the IBM Spectrum Control GUI

To configure IBM Spectrum Control for your environment, access the GUI.

For information about how to start the GUI, see [Opening IBM Spectrum Control GUIs and CLIs](#).

To set up IBM Spectrum Control for monitoring and managing resources, complete the following tasks:

Add resources for monitoring.

You can add the following resources for monitoring:

- Storage systems
- Servers
- Hypervisors
- Switches
- Fabrics
- NetApp Filers

Configure the retention of data in the database repository

Determine how long that IBM Spectrum Control stores data about the resources that are being monitored.

(Optional) Configure the collection of data about the sizes of GPFS snapshots

By default, IBM Spectrum Control does not collect information about the size of GPFS snapshots. To view that information, you must enable IBM Spectrum Control to collect information about the size of GPFS snapshots.

For information about how to configure collection of snapshot size information, see [Collecting information about the sizes of snapshots in IBM Spectrum Scale](#).

Configure alert notifications for conditions that are detected on monitored resources

Determining when and how you're alerted to configuration and performance conditions within your storage environment is important to helping you maintain and administer storage resources. Many conditions can trigger alerts. IBM Spectrum Control provides default alert policies for each resource type, or you can configure alert definitions so that IBM Spectrum Control examines the data about your resources for the conditions that you specify.

Specify the users that can access the product and the functions that are available to them

Assign roles to groups of users to determine which functions are available in IBM Spectrum Control. When a user ID is authenticated to IBM Spectrum Control through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

Deploy Storage Resource agents to set up data features

Deploy Storage Resource agents on servers to enable data collection and to monitor the fabrics that are visible to the server.

Monitor and manage resources

Access the IBM Spectrum Control GUI to monitor and manage your storage, including the following tasks:

- Monitor the status of resources
- View capacity and asset information about resources
- Troubleshoot the performance of resources
- Set performance thresholds and alerts
- View the relationships between resources
- Optimize storage
- Configure for the cloud and storage provisioning
- Generate detailed reports about resources

Remove the IBM Spectrum Control installation directory from the antivirus scan

Scans that are conducted by antivirus software cause more I/O write operations on the IBM Spectrum Control and server disk drives. For example, the antivirus software can be set to scan any file, record, or process that is written to or read from. The scan includes logs, read/write operations for database records, and Java™ processes that are used by IBM Spectrum Control. The antivirus software performs write operations on the objects that it identifies with these scans. These additional operations affect product performance and are visible in the performance monitoring tools for the operating system on which the IBM Spectrum Control server is installed. For example, on the Windows operating system, the Reliability and Performance Monitor might show an unusually high number of active write operations for IBM Spectrum Control. To improve performance, update your antivirus software so that it excludes the IBM Spectrum Control installation directory from scans.

McAfee tip for Windows: If McAfee Adaptive Threat Protection is enabled on the server where IBM Spectrum Control is installed, it might prevent some services from starting or stopping. To help avoid this issue, open McAfee Adaptive Threat Protection and go to settings. In the Real Protect Scanning (Windows only) section, verify if Enable client-based scanning is selected. If so, select Low from the Sensitivity level list.

For more information about McAfee Adaptive Threat Protection, see the following links:

- [Adaptive Threat Protection — Options](#)
- [Overview of Adaptive Threat Protection](#)

For more information about these tasks and how to manage and monitor your resources with IBM Spectrum Control, see [Managing resources](#).

Uninstalling

You can uninstall IBM Spectrum® Control by using the uninstallation program or by using silent mode from the command line.

About this task

To uninstall IBM Spectrum Control, uninstall these components and related software:

- Uninstall any Storage Resource agents installed on your remote systems.

- If all IBM Spectrum Control components are installed on one server, uninstall all components.
- If the database repository is installed on a separate server, uninstall the IBM Spectrum Control components on the first server, and then uninstall the database repository on the second server.
- Uninstall IBM® Db2®.

Note: On the Windows operating system, before you uninstall IBM Spectrum Control, if you installed the Monitoring Agent service, stop the Monitoring Agent for Windows OS - Primary and Monitoring Agent for Windows OS - Watchdog services.

- [Uninstalling IBM Spectrum Control in a single-server environment](#)
In a single-server environment, you can uninstall IBM Spectrum Control by using the uninstallation program. The uninstallation process requires minimal user interaction.
- [Uninstalling IBM Spectrum Control in a multiple-server environment](#)
You can uninstall IBM Spectrum Control by using the uninstallation program. This uninstallation process requires minimal user interaction.
- [Uninstalling IBM Spectrum Control by using silent mode](#)
You can uninstall IBM Spectrum Control by using silent mode. This mode is useful if your system is running from a terminal that cannot display graphics.
- [Manually uninstalling IBM Spectrum Control components on AIX or Linux](#)
If the IBM Spectrum Control uninstallation process fails, you must manually uninstall the components.
- [Manually uninstalling IBM Spectrum Control components on Windows](#)
If the IBM Spectrum Control uninstallation process fails, you must manually uninstall the components.
- [Uninstalling Storage Resource agents](#)
You can uninstall locally or remotely installed Storage Resource agents.
- [Uninstalling Db2](#)
Uninstall IBM Db2 after you uninstall IBM Spectrum Control. You must use the Db2 uninstallation program to uninstall Db2.
- [Uninstalling IBM Cognos Analytics](#)
To uninstall IBM Cognos® Analytics, follow the procedures in the Cognos Analytics documentation.

Uninstalling IBM Spectrum Control in a single-server environment

In a single-server environment, you can uninstall IBM Spectrum® Control by using the uninstallation program. The uninstallation process requires minimal user interaction.

Procedure

To uninstall the software in a single-server environment:

1. Log on to the IBM Spectrum Control host server with the appropriate user privileges.
2. Start the IBM Spectrum Control uninstallation program for your operating system.
 - For the Windows operating system:
 - Click Start > Control Panel > Programs > Programs and Features.
 - Highlight IBM Spectrum Control and click Uninstall/Change.
 - For the AIX® or Linux® operating systems:
Run the following command from the root directory:


```
installation_dir/_uninst/uninstall
```

 where *installation_dir* is where IBM Spectrum Control is installed.
3. Review the message that indicates that all the installed IBM Spectrum Control components are uninstalled.
4. Click Uninstall and click Next.
You cannot cancel the uninstallation process after it starts.
5. Restart the system if needed.
For the Windows operating system, you can restart the system now or later. You must restart the system to reinstall IBM Spectrum Control.

For the AIX or Linux operating system, you *do not* have to restart the system.
6. If an error occurred during the uninstallation process, review the installation log files.

Related reference

- [Planning for IBM Spectrum Control authentication and authorization](#)
- [Reviewing the log files to resolve installation issues](#)

Uninstalling IBM Spectrum Control in a multiple-server environment

You can uninstall IBM Spectrum® Control by using the uninstallation program. This uninstallation process requires minimal user interaction.

About this task

For this procedure, the terms *Server A* and *Server B* denote the two servers. The IBM Spectrum Control database repository is installed on *Server A* and the IBM Spectrum Control servers are installed on *Server B*.

Procedure

To uninstall the software in a multiple-server environment, complete the following steps:

1. Complete these steps on *Server B*:
 - a. Log on to *Server B* with the appropriate user privileges to uninstall the IBM Spectrum Control components.
 - b. Start the IBM Spectrum Control uninstallation program.
For the Windows operating system, click Start > Control Panel > Programs > Programs and Features. Highlight IBM Spectrum Control and click Uninstall/Change.
For the AIX® or Linux® operating system, run the following command from the root directory:

```
installation_dir/_uninst/uninstall
```


where *installation_dir* is the location where IBM Spectrum Control is installed.
A window is displayed that indicates that all IBM Spectrum Control components installed on the system will be uninstalled.
 - c. Click Uninstall. A confirmation is displayed.
 - d. Click Next to start the uninstallation process.
Tip:
Remember the following for your operating system:
 - For the Windows operating system, when the uninstallation process is finished, a window is displayed so you can choose to start your system now or later. You must restart the system before you can reinstall IBM Spectrum Control.
 - For the AIX or Linux operating system, you do not have to restart your system.
2. Complete these steps for *Server A*:
 - a. Log on to *Server A* with the appropriate user privileges to uninstall the IBM Spectrum Control database repository.
 - b. Start the IBM Spectrum Control uninstallation program.
 - c. Follow the instructions to uninstall the database repository.
3. If an error occurred during the installation process, review the installation log files.

Related reference

- [Planning for IBM Spectrum Control authentication and authorization](#)
- [Reviewing the log files to resolve installation issues](#)

Uninstalling IBM Spectrum Control by using silent mode

You can uninstall IBM Spectrum® Control by using silent mode. This mode is useful if your system is running from a terminal that cannot display graphics.

About this task

In silent mode, the following server configurations are possible:

Tip: The terms *Server A* and *Server B* denote the two servers in a multiserver environment.

Option 1

All IBM Spectrum Control components are installed on one server (single-server environment).

Option 2

The IBM Spectrum Control database repository is installed on *Server A* and all other IBM Spectrum Control components are installed on *Server B*.

To uninstall IBM Spectrum Control by using silent mode, complete the following steps:

Procedure

1. Complete the following steps for option 1:
 - a. Log on to the IBM Spectrum Control system with the appropriate user privileges.
 - b. Start the uninstallation program.
 - For the Windows operating system, run the following in the C:\ directory:

```
installation_dir\_uninst\uninstall.bat -i silent
```
 - For the AIX® or Linux® operating systems, run the following command from the root directory:

```
installation_dir/_uninst/uninstall -i silent
```


where *installation_dir* is where IBM Spectrum Control is installed.
 - c. Restart the system, if needed.
 - For the Windows operating systems, restart the system after the uninstallation process is complete.
 - For the AIX or Linux operating systems, you do not have to restart your system.
2. Complete the following steps for option 2:
 - a. Log on to *Server B* with the appropriate user privileges to uninstall IBM Spectrum Control.
 - b. Start the uninstallation program.
 - For the Windows operating system, run the following from the C:\ directory:

```
installation_dir\_uninst\uninstall.bat -i silent
```
 - For the AIX or Linux operating system, run the following command from the root directory:

```
installation_dir/_uninst/uninstall -i silent
```
 - c. Restart the system, if needed.
 - For the Windows operating systems, restart the system after the uninstallation process is complete.
 - For the AIX or Linux operating systems, you do not have to restart your system.
 - d. Log on to *Server A* with the appropriate user privileges to uninstall the IBM Spectrum Control database repository.
 - e. Start the uninstallation program.

- For the Windows operating system, run the following command from the C:\ directory:

```
installation_dir\uninst\uninstall.bat -i silent
```

f. Restart the system, if needed.

- For the Windows operating systems, restart the system after the uninstallation process is complete.
- For the AIX or Linux operating systems, you do not have to restart your system.

Related tasks

- [Manually uninstalling IBM Spectrum Control components on Windows](#)
- [Manually uninstalling IBM Spectrum Control components on AIX or Linux](#)

Related reference

- [Planning for IBM Spectrum Control authentication and authorization](#)
- [Reviewing the log files to resolve installation issues](#)

Manually uninstalling IBM Spectrum Control components on AIX® or Linux®

If the IBM Spectrum® Control uninstallation process fails, you must manually uninstall the components.

Before you begin

This procedure assumes the following conditions:

- You installed IBM Spectrum Control in the default directory.
- You used default ports.
- You used the default IBM Spectrum Control database repository named TPCDB.

Procedure

To manually uninstall IBM Spectrum Control components, complete the following steps:

1. Run the following command to stop the Storage Resource agent:

```
kill -9 `ps -aef | grep /opt/IBM/TPC/agent | grep -v grep | awk '{ print $2 }'` 2> /dev/null
```

2. Run the following command to stop the Data server:

```
kill -9 `ps -aef | grep data | grep -v grep | awk '{ print $2 }'` 2> /dev/null
```

3. Run the following command to stop the Device server:

```
kill -9 `ps -aef | grep deviceServer | grep -v grep | awk '{ print $2 }'` 2> /dev/null
```

4. Run the following command to stop the Alert server:

```
kill -9 `ps -aef | grep alertServer | grep -v grep | awk '{ print $2 }'` 2> /dev/null
```

5. Run the following command to stop the Export server:

```
kill -9 `ps -aef | grep export | grep -v grep | awk '{ print $2 }'` 2> /dev/null
```

6. Run the following command to stop the GUI:

```
kill -9 `ps -aef | grep webServer | grep -v grep | awk '{ print $2 }'` 2> /dev/null
```

7. Run the following command to stop the data collector:

```
kill -9 `ps -aef | grep collector | grep -v grep | awk '{ print $2 }'` 2> /dev/null
```

8. Drop the database by completing the following steps:

- a. Source the **db2profile**.

- b. Enter **db2 drop db TPCDB**.

9. Go to the **/db2_admin_home/db2_instance_name/** directory and ensure that the database directory is deleted.

10. Delete the following files and folders:

- **rm -r `find /etc -name *tpcdsrv* -print`**
- **rm -r `find / -name *itsanm* -print`**
- **rm -r /etc/Tivoli/TSRM/registry**
- **rm -r /etc/Tivoli/TSRM/registryNA**
- **rm -r /etc/Tivoli/TSRM/lock**
- **rm -fr /opt/IBM/TPC**

Manually uninstalling IBM Spectrum Control components on Windows

If the IBM Spectrum® Control uninstallation process fails, you must manually uninstall the components.

Before you begin

This procedure assumes that the following conditions are met:

- You installed IBM Spectrum Control in the default directory.
- You used default ports.
- The Storage Resource agent service name is *Agent1*.
- You used the default IBM Spectrum Control database repository name *TPCDB*.

About this task

To manually uninstall IBM Spectrum Control components, complete the following steps in a Windows command window:

Procedure

1. To uninstall the Storage Resource agent, run the following command:

```
sc stop Agent1
sc delete Agent1
```

2. To remove the Storage Resource agent registry entries, run the following command:

```
reg delete HKLM\SOFTWARE\IBM\TPCAGENT\1
```

where 1 is the instance of the Agent.

If you do not have multiple agents, you can also run the following command:

```
reg delete HKEY_LOCAL_MACHINE\SOFTWARE\IBM\TPCAGENT
```

Important: If you have multiple agents and you run this command, all of the agents are deleted.

3. To uninstall the Data server, run the following commands:

```
sc stop TSRMsrvt
sc delete TSRMsrvt
```

4. To stop the Device server and data collector, run the following command:

```
cd C:\Program Files\IBM\TPC\scripts
stopTPCdevice.bat
```

Tip: If the script does not exist, run the following command to stop the Device server:

```
sc stop deviceServer
```

If the script does not exist, run the following command to stop the data collector:

```
sc stop datacollector1
```

5. To delete the Device server, run the following command:

```
sc delete deviceServer
```

6. To delete the data collector, run the following command:

```
sc delete datacollector1
```

7. Run the following commands to uninstall the Alert server:

```
sc stop alertServer
sc delete alertServer
```

8. Run the following commands to uninstall the Export server:

```
sc stop exportServer
sc delete exportServer
```

9. Run the following commands to uninstall the Web server:

```
sc stop webServer
sc delete webServer
```

10. Run the following commands to delete the **wlp** directories:

```
cd C:\Program Files\IBM\TPC
del wlp
```

11. Run the following command to remove the IBM Spectrum Control directories:

```
rmdir /s/q "C:\Program Files\IBM\TPC"
```

12. Restart your computer.

Uninstalling Storage Resource agents

You can uninstall locally or remotely installed Storage Resource agents.

About this task

To uninstall a Storage Resource agent, use the GUI to remove the server that the agent is deployed on.

- [Uninstalling Storage Resource agents manually](#)
You can manually uninstall Storage Resource agents.
- [Deleting Storage Resource agent registry entries after a failed installation or uninstallation](#)
If your IBM Spectrum Control installation or uninstallation fails, you must find and manually delete the Storage Resource agent registry entries that remain.
- [Removing servers](#)
You can remove servers that you no longer want to monitor with IBM Spectrum Control.

Uninstalling Storage Resource agents manually

You can manually uninstall Storage Resource agents.

Procedure

To uninstall a Storage Resource agent, complete the following steps:

1. If the Storage Resource agent is on the Virtual I/O system, complete the following steps:
 - a. Log in with the **padmin** user ID.
 - b. Run the following command to set up the AIX® environment:

```
oem_setup_env
```

2. Go to the directory where the agent is installed:

```
agent_install_location
```

where *agent_install_location* is where the Storage Resource agent is installed.

3. Run the uninstallation command:

Restriction: Do not delete the Storage Resource agent component from the IBM Spectrum® Control server by using the GUI or the command-line interface. This component must exist on the server to complete the uninstall process.

```
➤ bin/Agent -uninstall -force -serverName TPC_server_name -debug MAX
```

Where:

-uninstall

Uninstalls the agent.

-force

This optional parameter forces an uninstallation. If you use this parameter, do not provide the **serverName** parameter.

-serverName *TPC_server_name*

TPC_server_name is the Data server name as defined in IBM Spectrum Control. You can check the configuration file for the server name:

```
agent_install_location/config/Agent.config
```

where *agent_install_location* is where the agent is installed.

-debug MAX

This optional parameter is for debugging purposes. If you set the -debug parameter, then some files are not deleted.

Tip: If you run the uninstallation program from the bin directory, the bin directory is not deleted.

If you run the uninstallation program outside of the agent installation directory, then you must specify the full path.

If the uninstallation fails, you must look at the return codes. For more information about return codes, see [Return codes used by Storage Resource agent](#).

Deleting Storage Resource agent registry entries after a failed installation or uninstallation

If your IBM Spectrum® Control installation or uninstallation fails, you must find and manually delete the Storage Resource agent registry entries that remain.

Deleting registry files in the Windows operating system

To delete registry entries after a failed installation or uninstallation in the Windows operating system:

1. In Windows Registry Editor, navigate to **HKEY_LOCAL_MACHINE > SOFTWARE > IBM > TPCAGENT > NUMBER** or **HKEY_LOCAL_MACHINE > SOFTWARE > IBM > TPCAGENT**.
2. Delete this entry:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\IBM\TPCAGENT\1]  
"Home"="C:\Program Files\IBM\TPC"  
"Port"="9567"
```

3. If the Storage Resource agent is running as a daemon service, delete the service.
For example, `sc delete AgentX`, where AgentX is the name of the Storage Resource agent.
4. Delete the folder in which the agent was installed.

Display the agent name and folder by running **services.msc** and right-click the service to see the properties.

- For daemon agents, verify that the associated CAP file was deleted.
The CAP file is located in the `%ALLUSERSPROFILE%\Application Data\IBM\CAP` folder. The naming convention for this file is `srainstance_default.xml`.

Remember: After uninstalling a daemon agent, you must restart your system.

Deleting registry files on the UNIX or Linux operating systems

To delete registry entries after a failed installation or uninstallation on the UNIX or Linux® operating systems:

- Delete the Storage Resource agent entry that is associated with the IBM Spectrum Control installation on the local server.
The terms *Server A* and *Server B* are used in this procedure to denote the different servers.

When you install IBM Spectrum Control on Server A, a Storage Resource agent is also installed on Server A. When you deploy another Storage Resource agent from IBM Spectrum Control Server B to Server A, an entry is added to the file `/etc/Tivoli/TSRM/registryNA` on Server A. This entry is associated with the Storage Resource agent deployed from Server B.

If a failure occurs when you install or uninstall the Storage Resource agent from Server A, you must delete the entry from the `/etc/Tivoli/TSRM/registryNA` file that is associated with the agent on Server A. For example, you would delete these lines from the file:

```
Instance=1
Home=/opt/IBM/TPC
Port=9567
```

- If the agent was a daemon, verify that the process is running and stop it.
 - To list all the running processes, run the following command:

```
ps -ef
```

- To stop a process, run the following command:

```
kill -9 process_number
```

- For daemon agents, verify that the associated CAP File was deleted.
The CAP file is located in the `/opt/IBM/CAP` directory. The naming convention for this file is `srainstance_default.xml`.

Removing servers

You can remove servers that you no longer want to monitor with IBM Spectrum® Control.

About this task

You can use the GUI to remove servers. If a Storage Resource agent is deployed to the server, the agent is uninstalled.

When the server is removed, it is no longer monitored by IBM Spectrum Control. All the data that was collected about the server is removed from the database repository.

Tip: When you remove a server, it is only removed from IBM Spectrum Control. The server is not physically deleted from the storage environment.

Procedure

To remove a server, complete the following steps:

- In the menu bar, go to Servers, > Servers.
- On the Servers page, right-click the server where the agent is deployed and select Remove.
- Click Remove to confirm that you want to remove the server.

Uninstalling Db2

Uninstall IBM® Db2® after you uninstall IBM Spectrum® Control. You must use the Db2 uninstallation program to uninstall Db2.

About this task

For information about uninstalling Db2® 11.1 on the Windows operating system, go to [Uninstalling your DB2® database product \(Windows\)](#).

For information about uninstalling Db2 11.5 on the Windows operating system, go to [Uninstalling your DB2 database product \(Windows\)](#).

- [Uninstalling Db2 on AIX or Linux](#)
You can uninstall IBM® Db2® on UNIX or Linux® systems.

Uninstalling Db2 on AIX or Linux

You can uninstall IBM® Db2® on UNIX or Linux® systems.

About this task

To uninstall Db2® on AIX® and Linux, complete the following steps:

Procedure

1. As the Db2 instance user (typically db2inst1), drop the databases, and stop Db2.
 - a. If you are the user with root authority, run the **su - db2inst1** command to switch the user to the Db2 instance user. Otherwise, log in as the Db2 instance user.
 - b. Run the **db2 list db directory** command to list the databases that might not have been removed by the IBM Spectrum® Control uninstallation program.
 - c. For each database name listed by the command in the previous step, run the **database_name** command. For example:

```
Command:
$ db2 drop db tcpdb
Response:
DB20000I The DROP DATABASE command completed successfully.
$ db2 drop db ibmcdm
Response:
DB20000I The DROP DATABASE command completed successfully.
```

- d. Run the **db2stop force** command. The following messages are displayed:

```
08/10/2015 21:46:13 0 0 SQL1064N DB2STOP processing was successful.
SQL1064N DB2STOP processing was successful.
```

- e. Run **db2 terminate**. The following message is displayed:

```
DB20000I The TERMINATE command completed successfully.
```

- f. Run the **exit** command to close the su shell or log out.

On Linux run the following commands:

```
groupdel dasadm1
groupdel db2iadm1
groupdel db2fadm1
```

7. Remove the Db2 directory.
 - a. Run the **cd** command.
 - b. Run the command to remove the Db2 directory on your operating system:
On AIX, run the **rm -r /opt/IBM/db2** command.

On Linux, run **rm -r /opt/ibm/db2** command.
8. Remove any Db2 definitions from `/etc/services`.

Uninstalling IBM Cognos Analytics

To uninstall IBM® Cognos® Analytics, follow the procedures in the Cognos Analytics documentation.

For information on uninstalling Cognos Analytics in a Microsoft Windows environment, go to [Uninstall IBM Cognos Analytics on Microsoft Windows operating systems](#).

For information on uninstalling Cognos Analytics in an UNIX or Linux® environment, go to [Uninstall IBM Cognos Analytics on UNIX or Linux operating systems](#).

Related information

- [Silent installation, uninstallation, and configuration](#)

Upgrading and migrating


Learn about the key tasks in the upgrade process for IBM Spectrum® Control.

The IBM Spectrum Control upgrade process guides you through the steps for upgrading to the current version.

Important: Upgrading the operating system of a server where IBM Spectrum Control is installed is not supported. If you upgrade the operating system of the server, IBM Spectrum Control might be unable to continue monitoring your resources. If you must upgrade the server, contact the IBM® Support team for IBM Spectrum Control Support to discuss your options.

Use the following image to find information about upgrading IBM Spectrum Control.

Tip: Click a step to view more information.



- [Preparing to upgrade](#)
You can use the following information to help prepare your upgrade of IBM Spectrum Control. Good planning helps you avoid delays and problems when you upgrade your system.
- [Backups](#)
Before you upgrade, you must back up the entire IBM Spectrum Control system.
- [Example of upgrading Db2 in a Windows environment](#)
To complete an upgrade from IBM Db2 11.1 to IBM Db2 11.5.7, stop the IBM Spectrum Control servers before you begin the upgrade procedure.
- [Example of upgrading Db2 in a AIX or Linux environment](#)
To complete an upgrade from IBM Db2 11.1 to IBM Db2 11.5.7, stop the IBM Spectrum Control servers before you begin the upgrade.
- [Upgrading IBM Spectrum Control in a single-server environment](#)
You can upgrade IBM Spectrum Control by using the installation program or silent mode.

- [Upgrading IBM Spectrum Control in a multiple-server environment](#)
To upgrade IBM Spectrum Control in a multiple-server environment, start the installation program on the server on which the database repository is installed. After you complete the installation program on this server, you can start the start the installation program on the other server.
- [Upgrading Storage Resource agents](#)
Upgrade Storage Resource agents to ensure that they are at the same release level as the IBM Spectrum Control server.
- [Upgrading SMI-S providers for storage systems](#)
Before you upgrade an SMI-S provider (also called the CIM agent or CIMOM agent), ensure that IBM Spectrum Control supports the SMI-S provider that you want and that the provider is compatible with the firmware versions of your storage systems.
- [Migrating alert definitions to alert policies](#)
When you upgrade to IBM Spectrum Control 5.3.2 or later, all of your monitored resources retain their existing alert definitions and are not included in an alert policy. However, at any time, you can place resources into an alert policy or create a policy based on a resource.
- [Scenario: Changing from LDAP and Local OS authentication to only LDAP or Local OS authentication after an upgrade](#)
In this scenario, you change from LDAP and Local OS authentication to only LDAP or Local OS authentication.

Preparing to upgrade

You can use the following information to help prepare your upgrade of IBM Spectrum® Control. Good planning helps you avoid delays and problems when you upgrade your system.

General considerations

Attention:

- You cannot upgrade from any release before IBM Spectrum Control 5.3.0. If you are currently running a version of IBM Spectrum Control earlier than 5.3.0, you must upgrade to IBM Spectrum Control 5.3.0 or later. For more information, see [IBM Spectrum Control interoperability matrix](#). When you are at IBM Spectrum Control 5.3.0 or later, you can upgrade to the highest version of IBM Spectrum Control.
- To help streamline, improve and clarify the product licenses for IBM Spectrum Control, the following licenses were renamed:
 - IBM Spectrum Control Standard Edition is now IBM Spectrum Control
 - IBM Spectrum Control Standard Select Edition is now IBM Spectrum Control Select Edition
 If you have the Advanced edition from previous releases, you are not affected by the renaming of product licenses. After you upgrade to 5.4.0, you can pick up right where you left off. However, keep in mind that all licenses are impacted by the deprecated features that are listed at [Discontinued features in IBM Spectrum Control](#).
- Upgrading the operating system of a server where IBM Spectrum Control is installed is not supported. If you upgrade the operating system of the server, IBM Spectrum Control might be unable to continue monitoring your resources. If you must upgrade the server, contact the IBM® Support team for IBM Spectrum Control Support to discuss your options.

Before you upgrade

Consider the following information before you upgrade:

- Before you upgrade IBM Spectrum Control, apply all available updates and patches to the operating system of the server where you installed IBM Spectrum Control. This is *not* the same as upgrading the operating system of the server where IBM Spectrum Control is installed, which is not supported.
- IBM Spectrum Control no longer supports Tivoli® Storage Productivity Center for Replication. To learn more about how to download and install IBM Copy Services Manager, see [IBM Copy Services Manager](#).
- The time and space considerations that are required to upgrade IBM Spectrum Control varies from site to site. Several factors, such as how many devices are being monitored and the version from which you are upgrading, affect the time that is required for the upgrade to complete. In some cases, the upgrade might take up to several hours to complete.
Note, that when you upgrade IBM Spectrum Control the database repository for metadata is reorganized to include all the latest changes and improvements in the new release. During the upgrade, the database reorganization process might take some time to complete or appear to be frozen. This behavior is typical, especially for larger environments or when you are upgrading from a much older release.

To help ensure that the upgrade and database reorganization complete successfully, make sure that you have enough space on your file system:

- On a Windows operating system: 1 GB - 2 GB on the drive where you installed IBM Spectrum Control.
- On a Linux® operating system: 100 MB - 300 MB in the /home file system and 1 GB - 2 GB in the / file system.
- On a AIX® operating system: 100 MB - 300 MB in the /home file system and 1 GB - 2 GB in the /opt file system.
- When you do upgrades of IBM Spectrum Control an offline reorganization of tables in the database repository occurs. These table reorganizations do not use temporary table space, so you need to consider the amount of free disk space you need to have before your upgrade. During an upgrade of IBM Spectrum Control, you must have free disk space equivalent to the largest amount of disk space that is being used by a single table in the database repository.

To determine the largest amount of disk space that is being used by a single table in the database repository, complete the following tasks on the computer where the IBM Spectrum Control database repository is installed:

1. Open a Db2® command prompt or switch to the Db2 instance owner.
2. Run the following commands:

```
a. db2 connect to <database name>

b. db2 "select MAX(physical_space) as needed_space from
(select
decimal((DATA_OBJECT_P_SIZE+INDEX_OBJECT_P_SIZE+LONG_OBJECT_P_SIZE+LOB_OBJECT_P_SIZE+XML_OBJECT_P_SIZE+COL_OBJEC
T_P_SIZE)/1024,10,2)
as physical_space from sysibmadm.admintabinfo where TABSCHEMA='TPC' order by physical_space desc)"

c. db2 disconnect <database name>
```

where <database name> is the name of the IBM Spectrum Control database repository, which is *TPCDB* by default.

Record the result from Step 2b and increase that value by 10 percent; this is the required amount of free disk space in megabytes.

The required free disk space must be available in each file system or drive where the database repository tables are stored.

To determine the file systems or drives where the database repository tables are stored, do the following tasks on the computer where the IBM Spectrum Control database repository is installed:

1. Examine the `installation_dir/config/InstallVariable.properties` file.
2. Record the value of the following parameters:

DFTDBPATH
CSTDBPATH

Each directory that is listed in the previously mentioned values is in a file system (AIX /Linux) or in a drive (Windows). The required free disk space must be available in each file system or drive.

During the upgrade of IBM Spectrum Control, the reorganization of tables in the database repository is logged to the `installation_dir/logs/log/traceInstaller_0.log` file.

- On Windows operating systems, ensure that all command prompts are closed. If you have open command prompts that have a current working directory in the IBM Spectrum Control installation directory structure, the upgrade might fail in various places because the upgrade process modifies certain directories.
- On Windows operating systems, ensure that the user who logs into the operating system to perform the upgrade of IBM Spectrum Control is an Administrator and does have the *Debug programs* privilege and does *not* have the *Deny access to this computer from the network* privilege in the Windows operating system security policy. These privileges can be found in Administrative Tools -> Local Security Policy -> Local Policies -> User Rights Assignment. After the upgrade of IBM Spectrum Control is complete, the required privileges for the user who performed the upgrade can be changed, but only if that user is *not* set as the Common User or the Db2 user in IBM Spectrum Control.
- The full X Window System must be installed to display the IBM Spectrum Control installation program on AIX and Linux operating systems.
- On Windows operating systems, if you installed the IBM Monitoring Agent, stop the Monitoring Agent for Windows OS Primary and Monitoring Agent for Windows OS Watchdog services. After your upgrade is complete, you can start these services again.

Security

Consider the following security information related to upgrades:

- Installing IBM Spectrum Control involves using your operating system in manners typical for installing new application software. If your antivirus software is set on the maximum mode, it might prevent some of your changes from being accepted. To verify that your upgrade completes correctly, enable your antivirus software product to allow the following instances:
 - The `/etc/hosts` file can be edited
 - Files can be created in the `/temp` directory
 - New executable files can be created in the `C:\Program Files` directory.McAfee tip for Windows: If McAfee Adaptive Threat Protection is enabled on the server where IBM Spectrum Control is installed, it might prevent some services from starting or stopping. To help avoid this issue, open McAfee Adaptive Threat Protection and go to settings. In the Real Protect Scanning (Windows only) section, verify if Enable client-based scanning is selected. If so, select Low from the Sensitivity level list. For more information about McAfee Adaptive Threat Protection, see the following links:
 - [Adaptive Threat Protection — Options](#)
 - [Overview of Adaptive Threat Protection](#)
- If you upgrade from a release earlier than IBM Spectrum Control 5.3.6 to the current version of IBM Spectrum Control, the SSL certificates that are used by the Data server and Storage Resource agents are not updated automatically to the SHA-256 hash algorithm. If you want to use the SHA-256 hash algorithm after you upgrade to the current version of IBM Spectrum Control, you must perform the steps in [Replacing default SSL certificates for the Data server and Storage Resource agents with custom SSL certificates](#).
- There are new certificate requirements from the [CA/B Forum](#) that are strictly enforced for macOS Catalina users that might affect your ability to access the IBM Spectrum Control GUI. During an upgrade of IBM Spectrum Control, certificates self-signed by IBM Spectrum Control will be made automatically compliant. However, if one or more of your certificates are not self-signed by IBM Spectrum Control, see [IBM Spectrum Control and macOS Catalina \(10.15\) Increased Security Policies on SSL Certificates](#) and validate that your certificates are compliant.
- Starting with IBM Spectrum Control 5.3.6, all default self-signed certificates have an expiration date of 825 days based on the new certificate requirements from the CA/B Forum. By default, upgrades of IBM Spectrum Control always renew the expiration of all default self-signed certificates used by the product for another 825 days. However, if you have installed your own certificates, IBM Spectrum Control does not modify those certificates for you.

After you upgrade

Consider the following information after you upgrade:

- If your upgrade process fails and a component cannot be upgraded, the upgrade process ends and the status of the component is shown as failed. When you start the upgrade process again, the installation program resumes upgrading the failed component.
- If you do a version upgrade of Db2 (For example, 11.1 to 11.5.7), you must apply a valid license to Db2 following the upgrade.
- After you upgrade to IBM Spectrum Control, all trace settings are reset to the default values. If you changed the trace settings, you must reset the trace settings after an upgrade.

Cognos Analytics

Consider the following information if you use Cognos® Analytics 11:

- If you plan to upgrade your Cognos Analytics 11 environment, see [How to upgrade your version of Cognos Analytics](#). After you upgrade IBM Spectrum Control and upgrade your Cognos Analytics 11 environment, you can upgrade the pre-defined reports by importing the pre-defined reports package from the most updated version of IBM Spectrum Control into Cognos Analytics. See [Importing IBM Spectrum Control pre-defined reports package](#).
- After you complete the upgrade of your Cognos Analytics 11 environment on Linux or AIX, verify that the link from the 32-Bit Db2 library to the `Cognos_installation_dir/bin/` directory is still present. If it is not present, perform Step 3 in [Copying Db2 files and linking the Db2 library on Linux and AIX](#).
- When you complete the upgrade of your Cognos Analytics 11 environment, and you save the Cognos Configuration, if you encounter a failure in the Check password for JVM truststore step, and it generates the **CAM-CRP-1613** error in the Details, see [How to regenerate cryptographic keys in Cognos Analytics 11](#).
- The IBM Cognos BI Reports component is no longer part of the IBM Spectrum Control product and is uninstalled during the upgrade. It is recommended that you back up your entire IBM Spectrum Control system before you proceed with an upgrade.

After you upgrade IBM Spectrum Control and are using IBM Cognos Analytics, import the pre-defined reports package into Cognos Analytics. See [Importing IBM Spectrum Control pre-defined reports package](#).

Related concepts

- [Product licenses](#)
- [Backups](#)

Related tasks

- [Replacing the default SSL certificate for the Device, Alert, or Web server with a self-signed certificate](#)
- [Replacing the default SSL certificate for the Export server](#)
- [Licensing Db2](#)

Related reference

- [Editing the response file](#)
- [Parameters in the silent Upgrade.properties file](#)
- [Planning for installation](#)

Backups

Before you upgrade, you must back up the entire IBM Spectrum® Control system.

The process involves backing up the following parts:

- IBM Spectrum Control database by using Db2® backup utilities.
For more information about backing up your database, see [Backing up the database offline using the command line](#).
- The entire IBM Spectrum Control server by using software such as IBM Spectrum Protect.

Tip: The frequency of backups depends on the data protection policies of your company, the IBM Spectrum Control work load, and the size of the IBM Spectrum Control environment. Your backup plan must include documentation with the user names and passwords for Db2 and IBM Spectrum Control.

If you experience problems during the migration, use these backups to restore your original IBM Spectrum Control system and try to migrate again. For more information about restoring the IBM Spectrum Control database, see [Restoring the database](#).

Restriction:

Do not restore the IBM Spectrum Control database from one version of IBM Spectrum Control into another version. For example, do not restore a backup from IBM Spectrum Control 5.3.3 into 5.3.4, or any other version.

Example of upgrading Db2 in a Windows environment

To complete an upgrade from IBM® Db2® 11.1 to IBM Db2 11.5.7, stop the IBM Spectrum® Control servers before you begin the upgrade procedure.

About this task

You can upgrade to the latest version of Db2 that IBM Spectrum Control supports before you install the latest version of IBM Spectrum Control.

Important: The following steps are an example of one way to upgrade from Db2 11.1 to Db2® 11.5.7 on a Windows operating system.

Procedure

1. Stop all of the Windows services for the IBM Spectrum Control servers, including IBM Cognos® Analytics (if you installed it on your system).
2. In a Db2 command window, run the following commands to force all users or applications to disconnect from Db2:

```
db2 force application all
db2 terminate
```

3. In a Db2 command window, run the following command to ensure that no applications are accessing Db2:

```
db2 list applications
```

4. Create a directory to back up the database.
For example: `C:\downloads\db2_backup`
5. In a Db2 command window, run the following command:

```
db2 BACKUP DATABASE database to C:\downloads\db2_backup
```

where *database* is the name of the IBM Spectrum Control database (usually TPCDB). If the Cognos Analytics database is installed (usually CM), make a similar backup of that database.

6. Stop all Db2 Windows services and exit Db2 in the Windows System Tray.
7. In Windows Explorer, go to the following directory:

```
C:\Db2_11.5.7_download_directory\SERVER_DEC\image
```

Where *Db2_11.5.7_download_directory* is where you downloaded the Db2 installation files.

8. Double-click the setup.exe file.

9. On the Db2 Setup Launchpad navigation tree view, click Install a Product.
10. On the Install a Product page, in the Db2 11.5.7 Server Editions section, click Work with Existing.
11. In the Work with an Existing Db2 Copy window, select the Db2 copy for Db2 11.1.
12. Click Launch Db2 Setup wizard.
13. On the Warning window about upgrading Db2, click OK.
14. In the Db2 Setup wizard, complete the following steps:
 - a. On the Welcome page, click Next.
 - b. On the Software License Agreement page, review and accept the license agreement, and click Next.
 - c. On the Select the installation type page, select Typical and click Next.
 - d. On the Select the installation, response file creation, or both page, keep the default selection and click Next.
 - e. On the Installation folder page, click Next.
 - f. On the Set the Db2 copy name page, click Next.
 - g. On the Set user information for the default Db2 instance page, enter the correct user name and password values for the existing Db2 11.1 installation, and click Next.
 - h. On the Enable operating system security for Db2 objects page, ensure that the settings and values are correct for the existing Db2 11.1 installation, and click Next.
 - i. If the Warning window about existing group names is displayed, click OK.
 - j. On the Start copying files and create response file page, click Finish.
 - k. If Warning windows are displayed about a restart that is required to complete the setup, click OK.
 - l. On the Setup is complete page, click Finish.
 - m. In the Db2 First Steps page, click Do not create profile.
 - n. Close the Db2 First Steps window.
15. Apply a valid license to Db2 and remove the trial license.
16. In a Db2 command window, run the following command to upgrade the IBM Spectrum Control database to Db2 11.5.7:

```
db2 UPGRADE DATABASE database USER user_name USING password
```

where *database* is the name of the IBM Spectrum Control database (usually TPCDB), *user_name* is the user who owns the Db2 instance where the database is located (usually db2admin), and *password* is the password that is associated with that user name.

17. In a Db2 command window, run the following command to upgrade the Cognos Analytics (if installed) database to Db2 11.5.7:

```
db2 UPGRADE DATABASE database USER user_name USING password
```

Where *database* is the name of the Cognos Analytics database (usually CM), *user_name* is the user who owns the Db2 instance where the database is located (typically db2admin), and *password* is the password that is associated with that user name. You might encounter the following message when you upgrade the Cognos Analytics database to Db2 11.5.7:

```
SQL0954C Not enough storage is available in the application heap to process the statement.
```

18. Restart your computer and ensure that all of the Windows services for the IBM Spectrum Control servers, including Cognos Analytics are started.

Related tasks

- [Licensing Db2](#)

Related information

- [Messages - SQL0750 - SQL0999](#)
- [Applying fix packs in Db2 database environments](#)
- [Upgrading a Db2 server \(Windows\)](#)
- [Upgrading a Db2 server \(Linux and UNIX\)](#)

Example of upgrading Db2 in a AIX or Linux environment

To complete an upgrade from IBM® Db2® 11.1 to IBM Db2 11.5.7, stop the IBM Spectrum® Control servers before you begin the upgrade.

About this task

You can upgrade to the latest version of Db2 that IBM Spectrum Control supports before you install the latest version of IBM Spectrum Control.

Important: The following steps are an *example* of one way to upgrade from Db2 11.1 to Db2 11.5.7 in a AIX® or Linux® environment.

Procedure

1. Stop all IBM Spectrum Control servers, including IBM Cognos® Analytics (if you installed it on your system).
2. Switch user to the Db2 instance owner (usually db2inst1).
3. From a command shell, run the following commands to force all users or applications to disconnect from Db2:

```
db2 force application all
```

```
db2 terminate
```

4. From a command shell, run the following command to ensure that no applications are accessing Db2:

```
db2 list applications
```

5. Create a directory to backup the database.
For example: /tmp/db2_backup

6. From a command shell, run the following command:

```
db2 BACKUP DATABASE database to /tmp/db2_backup
```

where *database* is the name of the IBM Spectrum Control database (usually TPCDB). If the Cognos Analytics database is installed (usually CM), make a similar backup of that database.

7. Switch user back to root.

8. From a command shell, change the directory to the location of the Db2 11.5.7 db2setup installation program and enter the following command:

```
./db2setup
```

9. On the Db2 Setup Welcome page, click New Install.

10. On the Choose a Product page, select Db2 Version 11.5.7.0 Server Editions and click Next.

11. On the Configuration page:

- Enter an installation directory or accept the default.
- Select Typical as the installation type.
- IMPORTANT:** Deselect Create an instance.
- Review and agree to the IBM terms and click Next.

12. On the Response File and Summary page:

- Select Install Db2 Server Edition on this computer and save my settings in a response file.
- Enter a response file name or accept the default.
- Review the Summary and click Finish

13. On the Setup has completed successfully page, review the information, and click Finish.

14. Switch user to the Db2 instance owner (usually db2inst1).

15. From a command shell, run the following command:

```
db2stop
```

16. Switch user back to root.

17. From a command shell, run the following commands:

```
installation_dir_for_Db2_11.5.7/instance/db2iupgrade Db2_instance
```

where *installation_dir_for_Db2_11.5.7* is the installation location of Db2 11.5.7 and *Db2_instance* is the Db2 instance where the database is located, usually db2inst1.

18. Switch user to the Db2 instance owner (usually db2inst1).

19. From a command shell, run the following command:

```
db2start
```

20. Apply a valid license to Db2 and remove the trial license.

21. From a command shell, run the following command to upgrade the IBM Spectrum Control database from Db2 11.1 to Db2 11.5.7:

```
db2 UPGRADE DATABASE database USER user_name USING password
```

where *database* is the name of the IBM Spectrum Control database (usually TPCDB), *user_name* is the user who owns the Db2 instance where the database is located (usually db2inst1), and *password* is the password associated with that user name.

22. From a command shell, run the following command to upgrade the Cognos Analytics (if installed) database from Db2 11.1 to Db2 11.5.7:

```
db2 UPGRADE DATABASE database USER user_name USING password
```

where *database* is the name of the Cognos Analytics database (usually CM), *user_name* is the user who owns the Db2 instance where the database is located (usually db2inst1), and *password* is the password associated with that user name.



23. Switch user back to root.

24. Start all of the IBM Spectrum Control servers, including Cognos Analytics (if you installed it on your system).

Related tasks

- [Licensing Db2](#)

Related information

-  [Applying fix packs in Db2 database environments](#)
-  [Upgrading a Db2 server \(Linux and UNIX\)](#)

Upgrading IBM Spectrum Control in a single-server environment

You can upgrade IBM Spectrum® Control by using the installation program or silent mode.

- [Upgrading IBM Spectrum Control in a single-server Windows environment by using the installation program](#)**
You can upgrade IBM Spectrum Control by using the installation program.
- [Upgrading IBM Spectrum Control in a single-server AIX or Linux environment by using the installation program](#)**
You can upgrade IBM Spectrum Control as a root user or non-root user on an AIX® or Linux® operating system by using the IBM Spectrum Control installation program. To upgrade IBM Spectrum Control as a non-root user, you must have sudo privileges on the target AIX or Linux server.
- [Upgrading IBM Spectrum Control in a single-server environment by using silent mode](#)**
You can upgrade IBM Spectrum Control by using silent mode. This upgrade method is useful if your computer cannot display graphics.

- [Editing the upgrade response file](#)
When you upgrade using the silent mode, use the `silent_Upgrade.properties` file.

Upgrading IBM Spectrum Control in a single-server Windows environment by using the installation program

You can upgrade IBM Spectrum® Control by using the installation program.

Before you begin

Before you upgrade, you must back up the entire IBM Spectrum Control system.

Procedure

To upgrade IBM Spectrum Control in a single-server Windows environment, complete the following steps:

1. Log on to the IBM Spectrum Control computer with the appropriate user privileges.
2. Start the IBM Spectrum Control installation program.
 - a. Download the image into a directory.
 - b. Extract the image files.
 - c. Enter the following command: `cd image_directory\SC` and then enter `setup.bat`.
3. Follow the prompts in the installation program to upgrade IBM Spectrum Control.
4. After the upgrade is finished, review the message log to ensure that no errors occurred.

Related concepts

- [Backups](#)

Related reference

- [Preparing to upgrade](#)
- [Reviewing the log files to resolve installation issues](#)
- [Planning for installation](#)
- [Starting the installation programs](#)

Upgrading IBM Spectrum Control in a single-server AIX or Linux environment by using the installation program

You can upgrade IBM Spectrum® Control as a root user or non-root user on an AIX® or Linux® operating system by using the IBM Spectrum Control installation program. To upgrade IBM Spectrum Control as a non-root user, you must have sudo privileges on the target AIX or Linux server.

Before you begin

Before you upgrade, you must back up the entire IBM Spectrum Control system.

Ensure that X Window System is installed on the target AIX or Linux server. The X Window System is required to upgrade IBM Spectrum Control by using the installation program.

If you want to upgrade IBM Spectrum Control as a non-root user, have your system administrator complete the following tasks on the target server before you begin the upgrade:

- If you are upgrading IBM Spectrum Control on an AIX server, install sudo on the server. The default AIX distribution does not include sudo privileges. To get the sudo installation package, go to <http://www.ibm.com/systems/power/software/ aix/linux/toolbox/date.html>. Find `sudo-version` in the Package column, and click RPM to download the package.
- Ensure that the non-root user has sudo privileges for the following commands on the server:
 - **xauth** (Enables the non-root user to provide access to their X Window System display for the root user.)
 - **setup.bin** (Enables the non-root user to launch the IBM Spectrum Control installation program.)

You must enter the full path to the commands when you configure the sudo privileges. For example:

```
username      ALL=(root)      SETENV: /usr/bin/xauth, /home/username/Downloads/IBMSC/SC/setup.bin
```

- If the sudo environment on the server uses the `secure_path` option, ensure that the following directories are included in the `secure_path` value for the non-root user who is upgrading IBM Spectrum Control:
 - `DB2_instance_owner_home_directory/sqllib/bin` (Example: `/home/db2inst1/sqllib/bin`)
 - `DB2_instance_owner_home_directory/sqllib/adm` (Example: `/home/db2inst1/sqllib/adm`)
 - `DB2_instance_owner_home_directory/sqllib/misc` (Example: `/home/db2inst1/sqllib/misc`)
- Ensure that the following requirements are met to enable the upgrade of IBM Spectrum Control by using the installation program:
 - The non-root user is able to use the X Window System on the target server and the X Window System **DISPLAY** environment variable is set correctly for the not-root user.
 - The **DISPLAY** environment variable is preserved in the sudo environment.
 - The root user has access to the X Window System display that is owned by the non-root user.

These display requirements do not apply if you upgrade IBM Spectrum Control in silent mode.

- If you are upgrading IBM Spectrum Control on an AIX server, ensure that the **ODMDIR** environment variable is preserved in the sudo environment.

Procedure

1. Log on to the IBM Spectrum Control computer as the root user or as the non-root user who has sudo privileges.
2. Start the IBM Spectrum Control installation program.

- a. Create a directory. Enter:

```
mkdir /SpectrumControl
```

- b. Download the image into the **SpectrumControl** directory.
- c. Enter the following command to extract the image files:

```
tar -xzf filename.tar.gz
```

- d. Ensure that the user name that you plan to use as the IBM Spectrum Control common user is in the **system** group and the **db2iadm1** group.
If the user name is not in a group, you can add the user as shown in the following command example. In this example, the user name is **db2inst1**:

```
usermod -G system db2inst1
```

Tip: If the user is already a member of secondary groups, the **usermod** command removes the **db2inst1** user name from all the secondary groups except for the system group. To preserve the existing secondary group memberships, provide a comma separated list of secondary groups after the **-G** flag. For more information, enter **man usermod** at the command prompt.

- e. Go to the **/SpectrumControl/SC** directory:

```
cd /SpectrumControl/SC
```

- f. Set up your shell environment to point to the instance where the database repository is installed. Source the **db2profile** for the instance that you want. For example, if the Db2® instance is **db2inst1**, you can source the **db2profile** by entering:

```
./home/db2inst1/sqllib/db2profile
```

Remember: There is a space between **.** and **/home**.

- g. Start the installation program by running the **./setup.bin** or **sudo -E ./setup.bin** command from the same command shell that you used in the previous step.
3. Follow the prompts in the installation program to upgrade IBM Spectrum Control.
 4. After the upgrade is finished, review the message log to ensure that no errors occurred.

Related concepts

- [Backups](#)

Related reference

- [Preparing to upgrade](#)
- [Reviewing the log files to resolve installation issues](#)
- [Planning for installation](#)
- [Starting the installation programs](#)

Upgrading IBM Spectrum Control in a single-server environment by using silent mode

You can upgrade IBM Spectrum® Control by using silent mode. This upgrade method is useful if your computer cannot display graphics.

Before you begin

Before you upgrade, you must back up the entire IBM Spectrum Control system.

Procedure

To upgrade IBM Spectrum Control in a single-server environment by using silent mode, complete the following steps:

1. Log on to the IBM Spectrum Control computer with the appropriate user privileges.
2. Edit and save the appropriate response file.
 - Edit the **silent_upgrade.properties** file and set the following parameters:
 - **LICENSE_ACCEPTED=true**
 - **CHOSEN_INSTALL_TYPE="Upgrade"**
3. Run the silent mode installation program.
 - On Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```

where *language* can be one of the following values:

- Czech - cs
- English - en
- French - fr
- German - de
- Hungarian - hu

- Italian - it
- Japanese - ja
- Korean - ko
- Polish - pl
- Brazilian Portuguese - pt_BR
- Russian - ru
- Spanish - es
- Chinese (Simplified) - zh_CN
- Chinese (Traditional) - zh_TW

absolute_path_to_response_file is the absolute path to the response file. For example, the following command specifies the language and the path:

```
setup.bat -l fr -i silent -f
c:\SpectrumControl\silent_Upgrade.properties
```

- On the AIX® or Linux® operating system, you must source the Db2® profile, **db2profile**, for the instance owner of the Db2 database. The following text is an example of the command you run to source the user profile:

```
./home/db2inst1/sql1lib/db2profile
```

- On AIX or Linux operating systems, run the following command:

```
./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

Where **absolute_path_to_response_file** is the absolute path to the response file. For example, the following command specifies the language and the path:

```
./setup.bin -l de -i silent -f /SpectrumControl/silent_Upgrade.properties
```

If you do not specify a language or if the language is not specified correctly, the installation program uses the English language as the default.

- Optional: Monitor the progress of the upgrade.

- To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
installation_dir\logs\traceTPCInstall.log
```

- To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

```
installation_dir/logs/traceTPCInstall.log
```

Related concepts

- [Backups](#)

Related reference

- [Planning for IBM Spectrum Control authentication and authorization](#)
- [Editing the upgrade response file](#)
- [Reviewing the log files to resolve installation issues](#)

Editing the upgrade response file

When you upgrade using the silent mode, use the **silent_Upgrade.properties** file.

Parameters in the silent_Upgrade.properties file

The **silent_Upgrade.properties** file contains the following parameters:

LICENSE_ACCEPTED=false

Specifies whether you accept the IBM Spectrum® Control license agreement. The options are:

true

You accept all the terms and conditions of the IBM Spectrum Control license agreement.

false or any other value

You do not accept the IBM Spectrum Control license agreement. The installation program exits.

CHOSEN_INSTALL_TYPE="Upgrade"

Specifies the upgrade type.

Important: Do not change the value for this parameter.

varAlertKeystorePW=default

varDeviceKeystorePW=default

varWebKeystorePW=default

Specify the password for the Alert server, Device server, and Web server keystore files.

The default keystore password is *default*. If any of the specified keystore passwords are incorrect, the upgrade exits.

Related tasks

- [Upgrading IBM Spectrum Control in a single-server environment by using silent mode](#)

Related reference

- [Creating a keystore for an IBM Spectrum Control server](#)

Upgrading IBM Spectrum Control in a multiple-server environment

To upgrade IBM Spectrum® Control in a multiple-server environment, start the installation program on the server on which the database repository is installed. After you complete the installation program on this server, you can start the start the installation program on the other server.

- [Upgrading to IBM Spectrum Control 5.3.0 or later with a remote database by using the installation program](#)
You can upgrade to IBM Spectrum Control 5.3.0 or later with a remote database by using the installation program.
- [Upgrading IBM Spectrum Control with a remote database by using silent mode](#)
You can upgrade IBM Spectrum Control in a multiple-server environment with a remote database by using silent mode.

Upgrading to IBM Spectrum Control 5.3.0 or later with a remote database by using the installation program

You can upgrade to IBM Spectrum® Control 5.3.0 or later with a remote database by using the installation program.

Before you begin

Before you upgrade, you must back up the entire IBM Spectrum Control system.

About this task

For this procedure, the terms *Server A* and *Server B* denote the two servers. Server A has Db2® Server Edition and the IBM Spectrum Control database repository installed. Server B has all of the other IBM Spectrum Control components installed.

Procedure

To upgrade using the installation program, complete the following steps:

1. Complete the following steps on Server B:
 - a. Log on to Server B with the appropriate user privileges.
 - b. On Server B, stop all IBM Spectrum Control services.
2. Complete the following steps on Server A:
 - a. Log on to Server A with the appropriate user privileges.
 - b. If necessary, upgrade Db2 on Server A to a version that is supported by IBM Spectrum Control.
For more information about the supported Db2 versions, see <http://www.ibm.com/support/pages/node/6249361#DB>. For more information about upgrading Db2, see [Upgrading Db2 servers](#).
Before you start the IBM Spectrum Control installation program on AIX® or Linux® operating systems, you must source the user profile, `db2profile`, for the instance owner of the Db2 database. The following text is an example of the command to source the user profile:

```
. /home/db2inst1/sqllib/db2profile
```
 - c. Start the IBM Spectrum Control installation program on Server A.
 - d. Follow the prompts in the installation program to upgrade the database repository on Server A.
 - e. After the upgrade is finished on Server A, review the message log to ensure that no errors occurred.
3. Complete the following steps on Server B:
 - a. Log on to Server B with the appropriate user privileges.
 - b. Start the IBM Spectrum Control installation program on Server B.
 - c. Follow the prompts in the installation program to upgrade the IBM Spectrum Control components on Server B.
 - d. After the upgrade is finished, review the message log file on Server B to ensure that no errors occurred.

Related concepts

- [Backups](#)

Related reference

- [Preparing to upgrade](#)
- [Planning for IBM Spectrum Control authentication and authorization](#)
- [Reviewing the log files to resolve installation issues](#)

Upgrading IBM Spectrum Control with a remote database by using silent mode

You can upgrade IBM Spectrum® Control in a multiple-server environment with a remote database by using silent mode.

Before you begin

Before you upgrade, you must back up the entire IBM Spectrum Control system.

About this task

For this procedure, the terms *Server A* and *Server B* denote the two servers. Server A has Db2® and the IBM Spectrum Control database repository installed. Server B has all of the other IBM Spectrum Control components installed.

Procedure

To upgrade IBM Spectrum Control by using silent mode, complete the following steps:

1. Complete the following steps on Server B:
 - a. Log on to Server B with the appropriate user privileges.
 - b. On Server B, stop all of the IBM Spectrum Control services.
See [Stopping the IBM Spectrum Control servers by using the GUI](#).
2. Complete the following steps on Server A:
 - a. Log on to Server A with the appropriate user privileges.
 - b. If necessary, upgrade Db2 on Server A to a version that is supported by IBM Spectrum Control.
For more information about the supported Db2 versions, see <http://www.ibm.com/support/pages/node/6249361#DB>. For more information about upgrading Db2, see [Upgrading DB2® servers](#).
 - c. Edit and save the appropriate response file.
 - Edit the `silent_Upgrade.properties` file and set the following parameters:
 - `LICENSE_ACCEPTED=true`
 - `CHOSEN_INSTALL_TYPE="Upgrade"`
 - d. On AIX® or Linux® operating systems, source the user profile (`db2profile`) for the instance owner of the Db2 database.
For example:

```
. /home/db2inst1/sqllib/db2profile
```

- e. Run the silent mode installation program on Server A.
 - For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```

where *language* can be one of the following values:

- Czech - cs
- English - en
- French - fr
- German - de
- Hungarian - hu
- Italian - it
- Japanese - ja
- Korean - ko
- Polish - pl
- Brazilian Portuguese - pt_BR
- Russian - ru
- Spanish - es
- Chinese (Simplified) - zh_CN
- Chinese (Traditional) - zh_TW

`absolute_path_to_response_file` is the absolute path to the response file. For example, the following command specifies the language and the path:

```
setup.bat -l fr -i silent -f  
c:\SpectrumControl\silent_Upgrade.properties
```

- For AIX or Linux operating systems, run the following command:

```
./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

For example, the following command specifies the language and the path:

```
./setup.bin -l de -i silent -f /SpectrumControl/silent_Upgrade.properties
```

If you do not specify a language or if the language is not specified correctly, the installation program uses the English language as the default.

- f. Optional: Monitor the progress of the upgrade on Server A.
 - To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
installation_dir\logs\traceTPCInstall.log
```

- To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

```
installation_dir/logs/traceTPCInstall.log
```

3. Complete the following steps on Server B:
 - a. Log on to Server B with the appropriate user privileges.
 - b. Edit and save the appropriate response file on Server B.
 - Edit the `silent_Upgrade.properties` file and set the following parameters:
 - `LICENSE_ACCEPTED=true`
 - `CHOSEN_INSTALL_TYPE="Upgrade"`
 - c. Run the silent mode installation program on Server B.
 - For Windows operating systems, run the following command:

```
setup.bat -l language -i silent -f absolute_path_to_response_file
```

where *language* can be one of the following values:

- Czech - cs
- English - en
- French - fr
- German - de
- Hungarian - hu
- Italian - it
- Japanese - ja
- Korean - ko
- Polish - pl
- Brazilian Portuguese - pt_BR
- Russian - ru
- Spanish - es
- Chinese (Simplified) - zh_CN
- Chinese (Traditional) - zh_TW

absolute_path_to_response_file is the absolute path to the response file. For example, the following command specifies the language and the path:

```
setup.bat -l fr -i silent -f  
c:\SpectrumControl\silent_Upgrade.properties.properties
```

- For AIX or Linux operating systems, run the following command:

```
./setup.bin -l language -i silent -f /absolute_path_to_response_file
```

For example, the following command specifies the language and the path:

```
./setup.bin -l de -i silent -f /SpectrumControl/silent_Upgrade.properties
```

If you do not specify a language or if the language is not specified correctly, the installation program uses the English language as the default.

- d. Monitor the progress of the upgrade on Server B.

- To view the progress of the silent mode installation on Windows operating systems, check the trace log in the following path:

```
installation_dir\logs\traceTPCInstall.log
```

- To view the progress of the silent mode installation on AIX or Linux operating systems, check the trace log in the following path:

```
installation_dir/logs/traceTPCInstall.log
```

Related concepts

- [Backups](#)

Related reference

- [Preparing to upgrade](#)
- [Editing the upgrade response file](#)
- [Planning for IBM Spectrum Control authentication and authorization](#)
- [Reviewing the log files to resolve installation issues](#)

Upgrading Storage Resource agents

Upgrade Storage Resource agents to ensure that they are at the same release level as the IBM Spectrum® Control server.

About this task

When you apply maintenance to IBM Spectrum Control, you can upgrade Storage Resource agents immediately or at a later time. To ensure that all your agents are at the current release level and to manage your network load, schedule upgrades regularly.

If a Storage Resource agent is not at the same level as the IBM Spectrum Control server, the following limitations occur:

- New functions in the current release might not be available for the resources that are monitored by the agent.
- Problem fixes are not applied to the agent.

You can upgrade Storage Resource agents by using the following methods:

- Use the Modify Agents > Upgrade action on the Servers page in the GUI.
- Use a Storage Resource agent command.

Procedure

To determine if a Storage Resource agent must be upgraded, complete the following steps:

1. In the menu bar, go to Servers > Servers.
2. View the values in the Agent State column.
If the state of the agent is Upgraded needed, the Storage Resource agent for the related server must be upgraded.

- [Starting agent upgrades](#)
Upgrade a Storage Resource agent to the same release level as the IBM Spectrum Control server.
- [Scheduling agent upgrades](#)
Schedule the upgrade process for a Storage Resource agent.
- [Upgrading Storage Resource agents by using a command](#)
You can manually upgrade Storage Resource agents.

Starting agent upgrades

Upgrade a Storage Resource agent to the same release level as the IBM Spectrum® Control server.

About this task

The ability to start the upgrade process for a Storage Resource agent is available when the following conditions are met:

- A Storage Resource agent must be deployed on the server that you want IBM Spectrum Control to monitor.
- An agent upgrade is not currently running for the server.
- The version of the agent that is deployed on the server is earlier than the IBM Spectrum Control server version.

Procedure

To upgrade a Storage Resource agent that was not upgraded at maintenance time, complete the following steps:

1. In the menu bar, go to Servers, > Servers.
2. On the Servers page, right-click the server that contains the Storage Resource agent to upgrade and select Modify Agents, > Upgrade.
3. Select Immediate from the Agent Upgrade list on the Upgrade Agent window.
4. Click Upgrade to start the upgrade process.

Scheduling agent upgrades

Schedule the upgrade process for a Storage Resource agent.

About this task

You can schedule the upgrade process for a Storage Resource agent when the following conditions are met:

- A Storage Resource agent must be deployed on the server that you want IBM Spectrum® Control to monitor.
- An agent upgrade is not currently running for the server.
- The version of the agent that is deployed on the server is earlier than the IBM Spectrum Control server version.

Tips:

- To manage the workload for a server and the network, schedule the agent upgrade for a time when the server and network are not busy.
- The scheduled time for an agent upgrade is based on the time zone of the IBM Spectrum Control server, not the time zone of the server where the Storage Resource agent is installed. For example, if an agent is installed on a server in the Central (CST) time zone, but the IBM Spectrum Control server is in the Pacific (PST) time zone, the time that is shown in the GUI when you schedule the upgrade is PST.

Procedure

To schedule the upgrade of a Storage Resource agent that was not upgraded at maintenance time, complete the following steps:

1. In the menu bar, go to Servers, > Servers.
2. On the Servers page, right-click the server that contains the Storage Resource agent to upgrade and select Modify Agents, > Upgrade.
3. Select Scheduled from the Agent Upgrade list on the Upgrade Agent window.
4. Select the date and time and click Upgrade to schedule the agent upgrade.

Upgrading Storage Resource agents by using a command

You can manually upgrade Storage Resource agents.

About this task

To manually upgrade the Storage Resource agent, complete the following steps:

Procedure

1. Go to the DVD location of the installation program (by using the **Storage Resource Agent** image) and go to the **bin** directory:

```
cd DVD_image_location/data/sra/operating_system_name
```

Where *DVD_image_location* is the location of the installation image for the Storage Resource agent.

2. Run the upgrade command:

```
bin/Agent -upgrade -installLoc agent_install_directory ->
- commType Daemon ->
1
```

Notes:

- 1 Parameter when the agent is run as a daemon service.

The parameters are:

-installLoc "*agent_install_directory*"

Location where the agent is installed. Enclose the directory name in quotation marks, for example, "C:\Program Files\IBM\TPC_SRA\".

-commType Daemon

If the agent is run as a daemon service, then this parameter must be specified.

Here is an example for a daemon-based service by using the default location:

```
bin/Agent -upgrade
-installLoc "/opt/IBM/TPC/"
-commType Daemon
```

Here is an example for a non-daemon service by using a non-default location:

```
bin/Agent
-upgrade -installLoc "C:\Program Files\IBM\TPC_SRA\"
```

Tip: If you run the upgrade program outside of the *DVD_image_location* installation directory, then you must specify the full path.

If the upgrade fails, see the return codes in [Return codes used by Storage Resource agent](#).

Upgrading SMI-S providers for storage systems

Before you upgrade an SMI-S provider (also called the CIM agent or CIMOM agent), ensure that IBM Spectrum® Control supports the SMI-S provider that you want and that the provider is compatible with the firmware versions of your storage systems.

To ensure that an SMI-S provider is supported, review the supported products list for the current release of IBM Spectrum Control. For more information about the products that are supported by IBM Spectrum Control 5.3.0 or later, see <https://www.ibm.com/support/pages/node/388393>. Under Storage, click the link for the appropriate release. If you are uncertain about the SMI-S provider support, contact IBM® customer support to help you with your upgrade plans.

- [Upgrading CIM agents](#)

To upgrade a Common Information Model (CIM) agent, you must have the upgrade instructions that are supplied by the provider of the agent.

Upgrading CIM agents

To upgrade a Common Information Model (CIM) agent, you must have the upgrade instructions that are supplied by the provider of the agent.

Procedure

To upgrade the CIM agent, follow these steps:

1. Stop all IBM Spectrum® Control server activity that might be dependent on the CIM agent (for example, performance monitor jobs, and discovery or probe jobs).
2. Stop the IBM Spectrum Control server services.
3. If any IBM Spectrum Control server processes remain, shut down the IBM Spectrum Control server for the duration of the CIM agent upgrade. If it is not possible to shut down the IBM Spectrum Control server, stop or kill the running processes.
4. Upgrade the CIM agent.
 - a. Refer to the instructions supplied by the CIM agent to perform the upgrade.
 - b. Verify that the CIM agent configuration is intact after the upgrade. Make sure that user accounts and passwords used with IBM Spectrum Control are still in place, and that all devices managed by the CIM agent are still listed.
5. Restart the IBM Spectrum Control server and IBM Spectrum Control server services.
6. Open the IBM Spectrum Control GUI.
7. Run a CIM object manager (CIMOM) discovery job.
8. Run a probe job for each storage system that is managed by the CIM agent.

Migrating alert definitions to alert policies

When you upgrade to IBM Spectrum® Control 5.3.2 or later, all of your monitored resources retain their existing alert definitions and are not included in an alert policy. However, at any time, you can place resources into an alert policy or create a policy based on a resource.

Use alert policies to manage the alert definitions and notification settings that apply to different sets of resources. Here are some reasons to place resources in an alert policy:

- You want to manage alert conditions and notification settings for a group of resources of the same type. For example, if you have several SAN Volume Controller storage systems in your environment, you can create an alert policy so that the alert definitions are the same for all of the SAN Volume Controller systems. Note that alert policies manage one type of resource only.
- You want to use one alert policy for the storage systems in your test environment, and another for the storage systems in your production environment.

Tips:

- Alert policies manage one type of resource only.
- A resource can be managed by only one alert policy.
- If a resource is managed by a policy, the resource cannot have alert definitions and notification settings that are independent of the policy. The alert definitions and notification settings that apply to the resource come from the policy.

It is not a requirement for resources to be managed by an alert policy. You can define alert conditions and notification settings for individual resources.

You can also define alerts and notification settings for applications and general groups to manage alerts for groups of resource components such as volumes or pools.

Migrating alert definitions from resources to alert policies

If you want to create a new alert policy that uses the alert definitions from one of your resources, complete the following steps:

1. Go to the details page for the resource from which you want to create the policy. For example, to create an alert policy from a FlashSystem 9100, click [Storage > Block Storage Systems](#).
2. Right-click the resource, then click View Alert Definitions.
3. Click Create Policy from the Policy Actions menu.
4. Type a name for the policy.
5. Make sure that the Use alert definitions and notification settings from *resource_name* option is selected.
6. Select any other resources that you want to add to the policy from the table of resources. Alert definitions in the policy are automatically applied to the resources you select. You can also choose to add resources later.

Scenario: Changing from LDAP and Local OS authentication to only LDAP or Local OS authentication after an upgrade

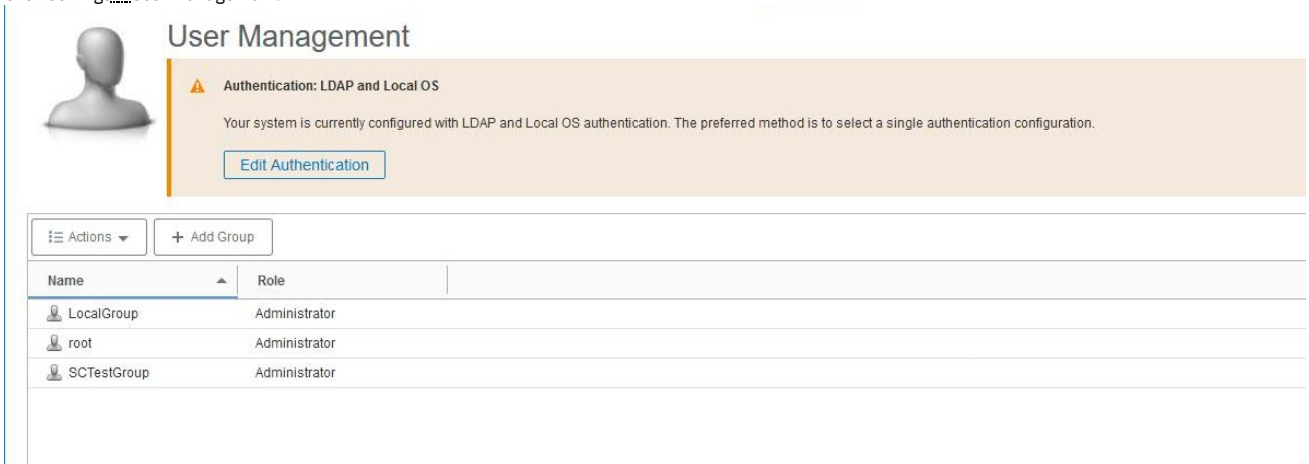
In this scenario, you change from LDAP and Local OS authentication to only LDAP or Local OS authentication.

About this task

You were running an earlier version of IBM Spectrum® Control with both LDAP and Local OS authentication configured. You upgraded to the latest version of IBM Spectrum Control. The preferred method of authentication is to select either LDAP or Local OS authentication.

Procedure

1. Back up the ldapregistry.xml file in the *installation_dir\wlp\usr\servers\webServer\registry* directory.
2. Log on as an administrator to the IBM Spectrum Control GUI.
3. Click [Settings > User Management](#).



4. Click Edit Authentication.
 - a. On the Authentication Configuration page, uncheck LDAP, if you want to change from LDAP and Local OS authentication to *only* Local OS authentication, and click Save.
 - b. On the Authentication Configuration page, uncheck Local OS, if you want to change from LDAP and Local OS authentication to *only* LDAP authentication and you do not want to alter your LDAP settings, and click Save.
 - c. On the Authentication Configuration page, uncheck Local OS, if you want to change from LDAP and Local OS authentication to *only* LDAP authentication and you do want to alter your LDAP settings, then start with Step 6 in [Changing from operating system to LDAP authentication](#)

Configuring

After IBM Spectrum Control is installed, you can configure it according to the standards and requirements of your storage environment.

- [Starting IBM Spectrum Control](#)
You can start IBM Spectrum Control by opening a web browser and entering a web address for the IBM Spectrum Control logon page. For example, you might enter

<https://storage.example.com:9569/srm>.

- [Configuring history and data retention](#)
Specify how long to retain the data that is collected about resources and the log files that are generated by IBM Spectrum Control. By specifying the number of weeks for history retention, you can control the amount of data that is retained and available for historical analysis and charting. The longer that you retain data, the more informative your analysis, but the more storage space that is required to store that data.
- [Configuring user authentication](#)
When IBM Spectrum Control is installed, default repositories are created, which allow you to control user access to the product.
- [Adding customized text to the logon page](#)
On the logon page for IBM Spectrum Control, you can show customized text when users access the GUI.
- [Managing a SAN without agents](#)
You can manage a SAN when there are no agents.
- [Setting timeout values for the Device server](#)
If a probe or discovery of a storage subsystem times out before the operation completes, you can increase the timeout values for the Device server.
- [Configuring Service Location Protocol](#)
You can enable IBM Spectrum Control to discover a larger set of storage devices through Service Location Protocol (SLP). In addition to some of the more common SLP configuration issues, there is also information about router configuration, SLP directory agent configuration, and environment configuration.
- [Configuring IP addressing](#)
This section provides information about configuring IP addressing.
- [Deploying Storage Resource agents](#)
You can manage your Storage Resource agent deployments.
- [Configuration guidelines for 500 or more agents](#)
You can use this information to help you manage 500 or more Storage Resource agents in IBM Spectrum Control.
- [Including a Storage Resource agent with a server golden image](#)
If you use a golden operating system image to deploy new servers in your environment, you can include the Storage Resource agent on that golden image. The golden image enables the agents to start and register with the IBM Spectrum Control server automatically upon deployment. This support applies only to Storage Resource agents running in daemon mode.
- [Checking for a fully qualified host name](#)
IBM Spectrum Control requires fully qualified host names. Some machines might be configured to return a short host name, such as system1 instead of a fully qualified host name, such as system1.tpc.example.com. This topic provides information on how to check for a fully qualified host name.
- [Granting local administrative privileges to a domain account](#)
Automatically grant administrative privileges to Windows domain accounts. The user account for the Storage Resource agent requires local administrative rights. Because these rights are not necessarily guaranteed for domain users in a Windows domain environment, you are shown how to grant local administrative rights to domain users. Using this procedure, you do not have to manually process each machine in the domain.
- [Importing authentication information for a Storage Resource agent](#)
The Storage Resource agent is installed as a non-daemon or daemon process. IBM Spectrum Control stores the authentication information to connect to the host on which the Storage Resource agent has installed for the non-daemon agent. This authentication information can be changed depending on the environment.
- [Installing and configuring the IBM Spectrum Control server with multiple NIC cards](#)
If your IBM Spectrum Control server has multiple network interface cards (NIC), install the IBM Spectrum Control server using a fully qualified hostname that resolves to the IP address of NIC card you want to use. After you install the server, all incoming and outgoing communication are successfully handled.
- [Replacing the default SSL certificate for the Device, Alert, or Web server with a self-signed certificate](#)
To replace the default SSL certificate for the Device, Alert, or Web server, with a self signed certificate, use the IBM® Key Management (iKeyman) utility.
- [Replacing the default SSL certificate for the Device, Alert, or Web server with a certificate from an external certificate authority](#)
To replace the default SSL certificate for the Device, Alert, or Web server, with a certificate from an external certificate authority, use the IBM Key Management (iKeyman) utility.
- [Replacing the default SSL certificate for the Export server](#)
You can replace the default SSL certificate for the Export server by adding the certificate and private key file to the appropriate directory and restarting the Export server.
- [Generating a new default self-signed SSL certificate for the Export server](#)
You can generate a new, default self-signed SSL certificate for the Export server by using the **openssl** command.
- [Enabling TLS 1.0 and 1.1 for ports](#)
IBM Spectrum Control uses Transport Layer Security (TLS) to secure communications between IBM Spectrum Control components.
- [Configuring Db2, AIX, and Linux for IPv6-only environment](#)
Use this information to configure Db2®, AIX®, and Linux® for an IPv6-only environment.

Starting IBM Spectrum Control

You can start IBM Spectrum® Control by opening a web browser and entering a web address for the IBM Spectrum Control logon page. For example, you might enter <https://storage.example.com:9569/srm>.

Before you begin

Before you start IBM Spectrum Control, ensure that you are using a supported web browser. For a list of web browsers that you can use with IBM Spectrum Control, see [IBM Spectrum Control - Platform Support: Servers, Agents, and Browsers - Web Browsers](#).

About this task

Start the IBM Spectrum Control GUI to administer and monitor the condition, capacity, and relationships of the resources within your storage environment.

Procedure

1. On a server running the Windows operating system, start IBM Spectrum Control GUI. For information about how to start the GUI on Windows operating systems, see [Opening IBM Spectrum Control GUIs and CLIs](#). If you are not on a server running the Windows operating system, start a web browser and enter the following address in the address field:

https://host_name:port/srm

In the preceding address, specify the following values:

host_name

The IBM Spectrum Control server. You can specify the host name as an IP address or a Domain Name System (DNS) name.

port

The port number for IBM Spectrum Control. The default port number for connecting to IBM Spectrum Control by using the HTTPS protocol is 9569. However, this port number might be different for your site. For example, the port number might be different if the default port range was not accepted during installation. If the default port number does not work, ask your IBM Spectrum Control administrator for the correct port number.

Tip: If you have a non-default port, check the value of the `WC_defaulthost_secure` property in `installation_dir/web/conf/portdef.props` file.

2. From the IBM Spectrum Control logon page, type your user name and password and click **Log in**.

The GUI opens in the browser.

Tip: If you want to log on to the GUI with Windows Domain credentials, use this form: `domain_name\user`.

Configuring history and data retention

Specify how long to retain the data that is collected about resources and the log files that are generated by IBM Spectrum® Control. By specifying the number of weeks for history retention, you can control the amount of data that is retained and available for historical analysis and charting. The longer that you retain data, the more informative your analysis, but the more storage space that is required to store that data.

About this task

Data that IBM Spectrum Control collects about a storage environment is stored in a DB2® database repository. The amount of data that is retained about resources can grow over time, and thus require more storage space for the repository. You can use the History Retention page to modify the data retention settings according to the monitoring and storage requirements of your environment. You must be assigned the Administrator role to modify data retention settings.

Procedure

1. In the menu bar, go to Settings > History Retention.
2. Click Edit to modify the following data retention settings:

Capacity history

Specify how long to retain a history of the capacity data that is collected about monitored resources. This value determines the amount of capacity data that is retained and available for historical analysis and charting. The longer that you retain data, the more informative your analysis, but the more storage space that is required.

Daily

Specify how long to retain capacity data that is collected daily about resources. You can retain daily data for up to 72 weeks and a minimum of 2 weeks.

Weekly

Specify how long to retain capacity data that is aggregated weekly for monitored resources. You can retain weekly aggregates for up to 96 weeks and a minimum of 4 weeks.

Monthly

Specify how long to retain capacity data that is aggregated monthly for monitored resources. You can retain monthly aggregates for up to 48 months and a minimum of 2 months.

Performance data

Specify how long to retain data that is collected by performance monitors.

Sample

Specify how long to retain sample data that is collected by performance monitors. Sample data represents the data that is collected each time a performance monitor is run. Because sample data is collected frequently, retaining that data can require significant disk space in the database repository. The required disk space is determined by the types of switches, storage systems, and number of volumes that are being monitored. You can retain sample data for up to 12 weeks.

Hourly

Specify how long to retain hourly data that is collected by performance monitors. You can retain hourly data for up to 24 weeks.

Daily

Specify how long to retain daily data that is collected by performance monitors. You can retain daily data for up to 156 weeks.

Consolidating performance data: Performance data is collected at intervals. An interval represents the number of minutes over which samples of performance data are averaged. When performance history is retained, data that is collected at certain intervals is automatically consolidated, or rolled up, to higher intervals. For example, data collected at 1-minute intervals is consolidated into 5-minute data; data collected at 5-minute intervals is consolidated into 1-hour data; and so on.

Tip: If performance data is collected at 1-minute intervals, the amount of data that is stored in the database repository increases significantly. The product stores only 7 days of sample data that is collected at 1-minute intervals.

Data for missing resources

Specify how long to retain data about internal resources that are no longer detected by IBM Spectrum Control. You can retain the data of removed resources for up to 52 weeks.

If the internal resource of a top-level resource is not detected when that top-level resource is probed, data about the resource is removed when the time limit is reached. The internal resource is removed only from the top-level resource that is probed. For example, if two weeks are specified, the data for a pool that is missing from a storage system for more than two weeks will be removed.

Only internal resources are automatically removed according to this setting. Storage systems, servers, hypervisors, switches, and fabrics must be removed manually.

Alert logs

Specify how long to retain alerts. An entry is generated each time that an alert condition is detected on a resource. Any alert that is older than this value is deleted. You can retain alerts for up to 12 weeks.

Job logs

Specify the maximum number of logs that are retained for data collection jobs. A log file is created each time that a job is run. When this number is reached, the entry for the oldest log is deleted. For example, if you accept the default value 5, and then run a probe 6 times, the log file for the first run is deleted. You can retain up to 20 logs for a job.

3. Click Save to apply the retention settings.
4. Optional: Click Restore Defaults to restore the retention settings to their default values.

Configuring user authentication

When IBM Spectrum® Control is installed, default repositories are created, which allow you to control user access to the product.

In the federated repositories framework, the following repositories are created:

File-based user repository

This repository contains the **tpcFileRegistryUser** user ID. This user password is the same as the Common User password that was entered during the IBM Spectrum Control installation.

When you use the password tools to change the IBM Spectrum Control passwords, the tpcFileRegistryUser user password gets changed so that it continues matching the Common User password.

If you encounter problems accessing IBM Spectrum Control using local operating system or LDAP credentials, you can log on to IBM Spectrum Control using the **tpcFileRegistryUser** user ID which is not affected by authentication configuration changes.

Operating system repository

In the federated repositories framework, the IBM Spectrum Control installation program creates two repositories on the IBM Spectrum Control web server. This server, which is in the installation_dir/wlp/usr/servers/webServer directory, is used as the primary WebSphere® Application Server Liberty server for user authentication in IBM Spectrum Control.

The Device server also runs on WebSphere Application Server Liberty, and it is only configured with the File-based user repository. If the web server is down, the Device server is used as the backup server to perform the user authentication and allows the common user name that was provided during IBM Spectrum Control installation and the **tpcFileRegistryUser** user ID to log on to IBM Spectrum Control.

You can add an LDAP repository and disable the operating system repository after you install IBM Spectrum Control; this configuration is completed in IBM Spectrum Control. The LDAP repository configuration settings are not propagated to the Device server. Therefore, if the web server is not running, the authorized LDAP users cannot log in to IBM Spectrum Control. The backup user authentication mechanism that is based on Device server allows the common user name that was entered during the IBM Spectrum Control installation, and the **tpcFileRegistryUser** user ID to be used to log on to IBM Spectrum Control.

If the computer is correctly configured with the Windows domain, the operating system repository also contains the domain users and groups that are managed by the Windows domain.

IBM Spectrum Control integrates with third party modules on the Linux® and AIX® operating systems for local user authentication. IBM Spectrum Control only supports the default module configuration settings with the AIX or Linux operating systems. The customization of configuration settings or using additional modules is not supported by IBM Spectrum Control.

The LDAP repositories that are supported by IBM Spectrum Control depend on WebSphere Application Server Liberty support. For more information, about the LDAP repositories that are supported, see <http://www.ibm.com/support/docview.wss?uid=swg27036471>.

WebSphere Application Server Liberty cannot resolve users or groups that are present in more than one repository in the federated repositories framework. Because of this limitation, you must select either the operating system repository or the LDAP repository for user authentication and authorization in IBM Spectrum Control. If you upgraded from an earlier version of IBM Spectrum Control with both the operating system repository and an LDAP repository configured, you can keep using both repositories. However, it is recommended that you select either the operating system repository or the LDAP repository.

The following table shows which user repositories are checked for IBM Spectrum Control authentication configurations when accessing IBM Spectrum Control using the IBM Spectrum Control GUI, the CLI, or the REST API:

Table 1. Authentication configurations and associated user repositories

Authentication configuration	User repositories checked
Local OS (default configuration)	Local operating system repository and file-based user repository
LDAP	LDAP repository and file-based user repository
LDAP and Local OS	Local operating system repository, LDAP repository, and file-based user repository

- **Authorizing users**
After IBM Spectrum Control is installed, you can assign roles to the user groups that are contained in the authentication repository. Roles determine the functions that are available to the users that are in a group.
- **Managing authentication**
The IBM Spectrum Control installation program establishes a default authentication configuration using the federated repositories feature of the WebSphere Application Server Liberty. You can configure and manage IBM Spectrum Control for LDAP authentication as a post-installation activity.

Authorizing users

After IBM Spectrum® Control is installed, you can assign roles to the user groups that are contained in the authentication repository. Roles determine the functions that are available to the users that are in a group.

About this task

The authentication repository can be an operating system repository or a Lightweight Directory Access Protocol (LDAP) repository. When IBM Spectrum Control is installed, the following user and groups are automatically configured for authentication to the product:

- User: tpcFileRegistryUser

- Windows group: Administrators group
- UNIX and Linux® group: root
- AIX® group: system

There are three IBM Spectrum Control roles that you can assign to user groups:

- Administrator (the Administrator, root, and system groups are automatically assigned to this role)
- Monitor
- External Application

Each role provides access to a specific set of functions. For more information about the functions that are available in each role, see [Role-based authorization](#).

- **[Role-based authorization](#)**
Roles determine the functions that are available to users of IBM Spectrum Control. When a user ID is authenticated to IBM Spectrum Control through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.
- **[Assigning a role to a group](#)**
Assign an IBM Spectrum Control role to one or more user groups. The role that is assigned to a group determines the product functions that are available to the users in that group.
- **[Determining the groups to which a user belongs](#)**
You can determine the groups to which a user belongs to and ensure that the user is in a group or groups that are assigned the correct IBM Spectrum Control role.
- **[Modifying the authentication mechanism](#)**
To modify how IBM Spectrum Control authenticates users and user groups, configure the authentication repository.
- **[Actions that are available based on role](#)**
Your IBM Spectrum Control role and product license determine the actions that are available in the product.

Role-based authorization


Roles determine the functions that are available to users of IBM Spectrum® Control. When a user ID is authenticated to IBM Spectrum Control through the GUI, CLI, or APIs, membership in an operating system or LDAP group determines the authorization level of the user.

The following table shows the IBM Spectrum Control roles and their authorization levels:

Table 1. IBM Spectrum Control roles and authorization levels

Roles	Authorization level
Administrator	<p>This role has full access to all monitoring and administrative functions. At least one group must have the Administrator role.</p> <p>Note: When IBM Spectrum Control is first installed, the following operating system groups are assigned the Administrator role:</p> <ul style="list-style-type: none"> • Windows: Administrators • UNIX and Linux®: root • AIX®: system
Monitor	<p>This role has access to the following read-only functions:</p> <ul style="list-style-type: none"> • Viewing and exporting information about monitored resources • Viewing, acknowledging, and removing alerts • Viewing tasks and data collection jobs • Opening management GUIs • Opening logs • Viewing chargeback, consumer, predefined capacity and inventory, and custom reports <p>Exception: Users with the Monitor role can provision storage if they are granted permission in a service class. A service class is a logical entity that describes storage capabilities and characteristics and can be used to specify requirements for storage provisioning. For more information about service classes, see Creating service classes.</p>
External Application	<p>If you assign the External Application role to the user, you must also assign one or more service classes to the user.</p> <p>This role does not enable users to log in to the IBM Spectrum Control GUI.</p>

Tips:

- To determine the role of the user who is logged in, click the user icon  in the upper-right corner of any page in the GUI.
- If a user belongs to multiple groups and the groups have different roles, the role with the highest level of authorization is granted to the user. For example, if a user belongs to a group that is assigned the Administrator role and also belongs to a group that is assigned a Monitor role, the user is granted the authorization of the Administrator role.
- If a user is not a member of a group that is assigned a IBM Spectrum Control role, no access is granted to that user.
- If assigned the Monitor role, a user can only open and view logs from the Data Collection page for the selected resource.

Nested groups are not supported: Adding active directory or any other type of domain user group to a local operating system group is not supported in IBM Spectrum Control. You can configure IBM Spectrum Control to authenticate domain IDs that rely on the operating system to perform the authentication operation against the active directory, but it cannot resolve nested groups.

Alternatively, you can configure LDAP authentication to perform queries against active directory user repositories and assign domain groups directly to roles within IBM Spectrum Control.

Assigning a role to a group

Assign an IBM Spectrum® Control role to one or more user groups. The role that is assigned to a group determines the product functions that are available to the users in that group.

About this task

If you are using LDAP authentication and you are using Microsoft Active Directory as your LDAP repository, do not assign the Active Directory Primary group for a user to an IBM Spectrum Control role. IBM Spectrum Control cannot identify user membership in a Primary group. Assign an IBM Spectrum Control role to a group that is not the Primary group for the user.

To assign a role to a group, complete the following steps:

Procedure

1. In the menu bar in the web-based GUI, go to Settings > User Management.
2. Click Add Group to search for groups that are defined in the authentication repository.
You can type the name of a group if you know its name, or specify a filter to search for existing groups in the authentication repository. For filters, use an asterisk (*) to represent unknown characters. You must enter at least one character in addition to an *.
For example, type `tpc*` to search for groups that begin with the letters "tpc" or "TPC". Type `*t` to search for groups that begin with or contain the letter "t" or "T".
3. In the list of groups, select one or more groups to which you want to assign a role.
4. In the Role field, select the role to assign to the group.
5. Click OK to assign the role.
The role that you select is applied to all the groups that you are adding. You can change the role assignments at any time after the group is added.

Related tasks

- [Determining the groups to which a user belongs](#)

Related reference

- [Role-based authorization](#)

Determining the groups to which a user belongs

You can determine the groups to which a user belongs to and ensure that the user is in a group or groups that are assigned the correct IBM Spectrum® Control role.

About this task

Use a command prompt to find the groups to which a user belongs.

Procedure

To determine the groups to which a user belongs, complete the following steps depending on your operating system:

1. Log on to the computer where the IBM Spectrum Control servers are installed, open a command prompt, and enter:

Option	Description
Windows standalone operating system	<code>net user <username></code>
Windows Domain	<code>net user <username> /DOMAIN</code>
Linux® and AIX® operating system	<code>groups <username></code>
LDAP	See your LDAP administrator for more information.

2. Verify that the user is in a group or groups that are assigned the correct IBM Spectrum Control role.

Modifying the authentication mechanism

To modify how IBM Spectrum® Control authenticates users and user groups, configure the authentication repository.

Before you begin

You must be assigned the Administrator role to modify the authentication repository and manage role and group assignments.

About this task

The authentication mechanism determines how IBM Spectrum Control authenticates users and the user groups that are available for assigned roles. During the installation process, the WebSphere® Application Server Liberty is configured with federated repositories. By default, authentication is configured to use the federated repositories, which contain a file repository and an operating system repository. The operating system repository includes the operating system users and groups that are defined on the server where IBM Spectrum Control is installed. For a server that is a member of a Windows domain, the operating system repository also includes the users and groups that are defined in that domain.

Important: WebSphere Application Server Liberty cannot resolve users or groups that are present in more than one repository in the federated repositories framework. Because of this limitation, you must select either the operating system repository or the LDAP repository for user authentication and authorization in IBM Spectrum Control. If you upgraded from an earlier version of IBM Spectrum Control with both the operating system repository and an LDAP repository configured, you can keep using both repositories. However, it is recommended that you select either the operating system repository or the LDAP repository.

Procedure

To modify your authentication, complete the following steps:

1. In the menu bar, go to Settings > User Management.
2. On the User Management page, click Edit Authentication.
The Authentication Configuration page is displayed.
3. Make your modifications.

Related concepts

- [Configuring user authentication](#)

Related tasks

- [Changing from operating system to LDAP authentication](#)

Actions that are available based on role

Your IBM Spectrum® Control role and product license determine the actions that are available in the product.

Users who are assigned the Administrator role or the Monitor role can use product functions. The actions that are available for each function depend on the role that is assigned to the user:

Administrator role

Users who are assigned the Administrator role have access to all monitoring and administrative actions.

Monitor role

Users who are assigned the Monitor role can view information about monitored resources and other objects such as tasks, alerts, and service classes. They can acknowledge alerts and resource statuses, open logs, and open management GUIs.

The following table outlines the actions that are available only for the Administrators role. All other actions are available to the Monitor and Administrator roles. In addition to the restrictions listed in this table, users who are assigned the Monitor role do not have access to user management functions.

Table 1. Product actions that are available only to users with the Administrator role

Function	Actions that require the Administrator role
Single dashboard view of the storage environment that you can use to manage storage systems, hypervisors, servers, and Fibre Channel fabrics.	<ul style="list-style-type: none">• Adding and removing resources• Administering connections• Scheduling data collection• Changing and viewing the automated probe schedule• Viewing and editing history retention settings• Modifying license settings
Performance monitoring for storage systems and Fibre Channel networks.	<ul style="list-style-type: none">• Scheduling performance monitors• Starting or stopping performance monitors
Capacity and usage monitoring of resources.	<ul style="list-style-type: none">• Scheduling probes• Starting or stopping probes• Modifying Storage Resource agents• Enabling automatic zoning
Health and alerting for hypervisors, networks, servers, and storage systems.	<ul style="list-style-type: none">• Creating, modifying, and deleting alert policies• Setting which alert policy manages a resource• Adding and modifying resources for management by an alert policy• Defining and modifying alert definitions• Editing alert notification settings

Function	Actions that require the Administrator role
Capacity and performance of the storage that applications, departments, and general groups use.	<ul style="list-style-type: none"> • Creating applications, departments, and general groups • Creating, modifying, and removing filters to add resources to applications • Adding and removing resources in applications and general groups, directly • Adding applications as members of other applications • Adding departments to other departments • Adding applications to departments
Storage reclamation	Viewing volumes that can be reclaimed
Roll-up reporting, in which capacity data is combined from multiple instances of IBM Spectrum Control for reporting purposes.	<ul style="list-style-type: none"> • Adding and removing subordinate servers • Starting a probe for a subordinate server • Modifying the connection information for a subordinate server
Predefined Reports <ul style="list-style-type: none"> • Predefined capacity reports allow users to quickly create reports about capacity anomalies and shortfalls, which can be scheduled and sent by email or saved to the user's file system, or both. • Predefined inventory reports allow users to quickly create reports about their storage resources, which can be scheduled and sent by email or saved to the user's file system, or both. 	<ul style="list-style-type: none"> • Creating, deleting, and editing reports • Configuring the email server • Emailing reports • Saving reports to the file system
Custom reports From any table view in the GUI, custom reports can be created, which can be scheduled and sent by email or saved to the user's file system, or both, about capacity of storage resources, the configuration and attributes of storage resources, and the performance of storage resources.	<ul style="list-style-type: none"> • Creating, deleting, and editing reports • Configuring the email server • Emailing reports • Saving reports to the file system
Chargeback and consumer reports <ul style="list-style-type: none"> • Chargeback reports show the capacity and the cost of the storage that is used by applications, departments, hypervisors, and physical servers. • Consumer reports show the capacity and the cost of the block storage that is used by an application, department, hypervisor, and physical server. 	<ul style="list-style-type: none"> • Creating, deleting, and editing reports • Configuring the email server • Emailing reports
Capacity limits for block storage systems and pools If your company has a policy to set a limit on the capacity that is used, you can set a capacity limit. When the capacity limit is set, you can then monitor the amount of capacity that is available before the capacity limit is reached.	<ul style="list-style-type: none"> • Setting capacity limits • Defining alerts for capacity limits • Removing capacity limits

Managing authentication

The IBM Spectrum® Control installation program establishes a default authentication configuration using the federated repositories feature of the WebSphere® Application Server Liberty. You can configure and manage IBM Spectrum Control for LDAP authentication as a post-installation activity.

Important: WebSphere Application Server Liberty cannot resolve users or groups that are present in more than one repository in the federated repositories framework. Because of this limitation, you must select either the operating system repository or the LDAP repository for user authentication and authorization in IBM Spectrum Control. If you upgraded from an earlier version of IBM Spectrum Control with both the operating system repository and an LDAP repository configured, you can keep using both repositories. However, it is recommended that you select either the operating system repository or the LDAP repository.

- [Changing from operating system to LDAP authentication](#)
You can configure IBM Spectrum Control to communicate with an external LDAP repository. For example, IBM® Tivoli® Directory Server or Microsoft Active Directory. This makes IBM Spectrum Control available to a larger set of users and groups. You are able to log in to IBM Spectrum Control with one set of credentials.
- [Changing from LDAP to operating system authentication](#)
To change from LDAP authentication to operating system authentication in IBM Spectrum Control, you can use the IBM Spectrum Control GUI.
- [Configuring user authentication alternatives](#)
If you encounter any issues with the recommended procedures for configuring user authentication, see the following topics for support.
- [Using the `ldapEntityType` element for advanced LDAP configuration](#)
To narrow the IBM Spectrum Control view of your LDAP structure so that you can find and map your groups to IBM Spectrum Control roles, you need to configure the `ldapEntityType` element within the `ldapRegistry` element.

Changing from operating system to LDAP authentication

You can configure IBM Spectrum® Control to communicate with an external LDAP repository. For example, IBM® Tivoli® Directory Server or Microsoft Active Directory. This makes IBM Spectrum Control available to a larger set of users and groups. You are able to log in to IBM Spectrum Control with one set of credentials.

Before you begin

When you change the authentication configuration, IBM Spectrum Control is available to users and groups in other repositories.

Important: WebSphere® Application Server Liberty cannot resolve users or groups that are present in more than one repository in the federated repositories framework. Because of this limitation, you must select either the operating system repository or the LDAP repository for user authentication and authorization in IBM Spectrum Control. If you upgraded from an earlier version of IBM Spectrum Control with both the operating system repository and an LDAP repository that is configured, you can keep using both repositories. However, it is recommended that you select either the operating system repository or the LDAP repository.

Procedure

1. Back up the `ldapregistry.xml` file in the `installation_dir/wlp/usr/servers/webServer/registry/` directory.
2. Log on as an administrator to the IBM Spectrum Control GUI.
3. Click Settings > User Management.
4. Click Edit Authentication.
5. On the Authentication Configuration page, select LDAP.
6. Click Download Files.
7. Save and extract the `ldapExamples.zip` file to the computer where you run your browser.
8. Use the information to edit the XML template file for your vendor.

For example, if your LDAP server is IBM Tivoli Directory Server, edit the `IBMDirectoryServer.xml` file and if your LDAP server is Microsoft Active Directory, edit the `ActiveDirectoryServerDefault.xml` file.

Edit the following parameters:

`id`

The unique identifier for the LDAP repository, which identifies the repository in the realm, for example, `LDAP1`.

`host`

The hostname of the primary LDAP server. The hostname is either the IP address or the computer name in a domain name system (DNS).

`sslEnabled`

Indicates whether SSL is used to connect to the LDAP server.

Important: If you set this parameter to `true`, and set the `port` parameter to the LDAP server secure communications port, when you upload the edited XML template file, IBM Spectrum Control downloads the SSL certificate from the LDAP server. Then, it is added to the Web server keystore. You must restart the Web server.

`port`

The port number for the LDAP server. By default, the port number for secure communication is 636, and for non-secure communication is 389.

Tip: Depending on the configuration of your LDAP server, you can specify a different port number.

`baseDN`

The baseDN (Distinguished Name) is the starting point for searches for users in the LDAP directory server. For example, if you have a DN value of `cn=John Doe, ou=rochester, o=ibm, c=us`, you can specify the LDAP base entry as any of the following options:

- `ou=rochester, o=ibm, c=us`
- `o=ibm, c=us`
- `c=us`

Important: The DN value that you enter in this field must be extensive enough to include all of the groups to which the users belong. For example, if a user in `ou=rochester, o=ibm, c=us` is also a member of groups that are in `ou=stategroups, o=ibm, c=us`, enter `o=ibm, c=us`.

If you want to set multiple baseDN parameters for your LDAP authentication configuration, then you must create a separate `<ldapRegistry>` entry in the XML template file for each unique baseDN parameter. For example:

```
<server description="IBM Web Server">
  <ldapRegistry activatedFilters="active_dir_server1" baseDN="ou=Marketing,dc=storage,dc=ibm,dc=com"
  bindDN="cn=Administrator,cn=users,dc=storage,dc=ibm,dc=com" bindPassword="password"
  host="ldap.storage.ibm.com"
  id=" _home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_configuring_fqz0_t_config_ldap_active_directory_I
  DAP1" ignoreCase="true" ldapType="Microsoft Active Directory"
  port="389" realm="TPCRealm" sslEnabled="false"/>
  <activatedLdapFilterProperties groupFilter="(& (cn=%v) (objectcategory=group))"
  groupIdMap="*:cn" groupMemberIdMap="memberof:member"
  id=" _home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_configuring_fqz0_t_config_ldap_active_directory_a
  ctive_dir_server1" userFilter="(& (sAMAccountName=%v) (objectcategory=user))"
  userIdMap="user:sAMAccountName"/>
  <ldapRegistry activatedFilters="active_dir_server2" baseDN="ou=Sales,dc=storage,dc=ibm,dc=com"
  bindDN="cn=Administrator,cn=users,dc=storage,dc=ibm,dc=com" bindPassword="password"
  host="ldap.storage.ibm.com"
  id=" _home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_configuring_fqz0_t_config_ldap_active_directory_I
  DAP2" ignoreCase="true" ldapType="Microsoft Active Directory"
  port="389" realm="TPCRealm" sslEnabled="false"/>
  <activatedLdapFilterProperties groupFilter="(& (cn=%v) (objectcategory=group))"
  groupIdMap="*:cn" groupMemberIdMap="memberof:member"
  id=" _home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_configuring_fqz0_t_config_ldap_active_directory_a
  ctive_dir_server2" userFilter="(& (sAMAccountName=%v) (objectcategory=user))"
  userIdMap="user:sAMAccountName"/>
  <ldapRegistry activatedFilters="active_dir_server3" baseDN="ou=Management,dc=storage,dc=ibm,dc=com"
  bindDN="cn=Administrator,cn=users,dc=storage,dc=ibm,dc=com" bindPassword="password"
  host="ldap.storage.ibm.com"
  id=" _home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_configuring_fqz0_t_config_ldap_active_directory_I
```

```

DAP3" ignoreCase="true" ldapType="Microsoft Active Directory"
port="389" realm="TPCRealm" sslEnabled="false"/>
<activeLdapFilterProperties groupFilter="(& (cn=%v) (objectcategory=group))"
groupIdMap="*:cn" groupMemberIdMap="memberof:member"
id=" home markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_configuring_fqz0_t_config_ldap_active_directory_a
ctive_dir_server3" userFilter="(& (sAMAccountName=%v) (objectcategory=user))"
userIdMap="user:sAMAccountName"/>
<federatedRepository>
  <primaryRealm allowOpIfRepoDown="true" name="TPCRealm">
    <participatingBaseEntry name="ou=Marketing,dc=storage,dc=ibm,dc=com"/>
    <participatingBaseEntry name="ou=Sales,dc=storage,dc=ibm,dc=com"/>
    <participatingBaseEntry name="ou=Management,dc=storage,dc=ibm,dc=com"/>
    <!-- The next two entries must NOT be changed -->
    <participatingBaseEntry name="o=TPCRealm"/>
    <participatingBaseEntry name="o=OSRealm"/>
  </primaryRealm>
</federatedRepository>
</server>

```

The preceding example contains the following information:

- Each `<ldapRegistry>` entry contains the identical values for the host, port, sslEnabled, bindDN, and bindPassword parameters.
- Each `<ldapRegistry>` entry contains a unique value for the baseDN and id parameters.
- Each `<ldapRegistry>` entry references its' own unique `<activeLdapFilterProperties>` entry.
- The `<federatedRepository>` entry contains multiple `<participatingBaseEntry>` entries, with each one matching one of the baseDN values in the `<ldapRegistry>` sections.

bindDN

The distinguished name that WebSphere Application Server Liberty uses when it binds to the LDAP repository. If no name is specified, WebSphere Application Server Liberty binds anonymously to the LDAP repository. In most cases, the bindDN and bindPassword values are required. However, when an anonymous bind satisfies all of the required functions, the bindDN and bindPassword values can be left blank except for the double quotes, which must have no spaces between them.

For example:

```

bindDN=""
bindPassword=""

```

If you are not sure whether an anonymous bind satisfies the required functions, contact your LDAP server administrator.

Attention: No single value for the bindDN parameter is correct for every Active Directory Server or for every LDAP server. The correct value for the bindDN parameter depends on the configuration of your Active Directory Server or your LDAP server. If you are unsure about the correct value to use for the bindDN parameter, contact your LDAP server administrator.

If you are using Active Directory as your LDAP repository and you know the *Active_Directory_user's_samAccountName_value*, but you want the Active Directory user full distinguished name in order to use that value as the bindDN parameter, run the following command on the Active Directory machine:

```
dsquery user -samid Active_Directory_user's_samAccountName_value
```

Example:

```

C:\Users\Administrator>dsquery user -samid SCAdministratorMSAD
"CN=SCAdministratorMSAD,CN=Users,DC=vcloud101dc,DC=local"

```

For more information about the **dsquery** command, see <https://social.technet.microsoft.com/wiki/contents/articles/2195.active-directory-dsquery-commands.aspx?PageIndex=3>

bindPassword

The password that WebSphere Application Server Liberty uses when it binds to the LDAP repository.

If the bindPassword parameter is already encrypted in the XML file, enter only an LDAP user name and password to test the pending LDAP authentication configuration. After your LDAP credentials are validated, you must immediately map an LDAP group to an IBM Spectrum Control role in the GUI before you log out with your Local OS credentials.

participatingBaseEntry

You must set this value to the same value as you set for the baseDN parameter or the federation for the LDAP repository fails. For example, if you set the baseDN parameter to `ou=rochester, o=ibm, c=us`, you must set the participatingBaseEntry parameter to: `<participatingBaseEntry name="ou=rochester, o=ibm, c=us" />`.

Important: Do not change these participatingBaseEntry parameters in the XML template file:

```

<participatingBaseEntry name="o=TPCRealm" />
<participatingBaseEntry name="o=OSRealm" />

```

9. Save the XML template file.
10. On the Authentication Configuration page, click Browse.
11. On the File Upload page, select the XML template file that you previously edited and click Open.
The XML template file is then uploaded to the IBM Spectrum Control server.
12. After IBM Spectrum Control downloads the SSL certificate from the LDAP server and adds it to the Web server keystore, click Restart Web Server.
Note: While the Web server is restarting, do not refresh your browser or attempt to navigate to another part of the GUI.
13. When the Web server is back online, on the LDAP Settings page, enter your LDAP user name, password, group name and click Save.
Note: The user name must be a member of the group.
If your change from Local OS authentication to LDAP authentication is successful, you are logged out of the IBM Spectrum Control GUI. You can log into the GUI using your LDAP credentials. You cannot log into the GUI using your operating system credentials.

If your change from Local OS authentication to LDAP authentication is *not* successful, click Discard. After the confirmation, your previous Local OS authentication is restored and you are returned to the User Management page.

Tip:

If you encounter any issues with above mentioned procedure, see [Configuring user authentication alternatives](#).

Results

You have changed from local operating system authentication to LDAP authentication in IBM Spectrum Control. You can log on to the IBM Spectrum Control GUI with LDAP credentials.

Related tasks

- [Exporting SSL certificate from the IBM Security Directory Server to a file](#)
- [Adding the SSL certificate for the LDAP server to the web server keystore that uses the IKEYCMD command](#)

Related reference

- [Role-based authorization](#)

Changing from LDAP to operating system authentication

To change from LDAP authentication to operating system authentication in IBM Spectrum® Control, you can use the IBM Spectrum Control GUI.

Before you begin

WebSphere® Application Server Liberty cannot resolve users or groups that are present in more than one repository in the federated repositories framework. Because of this limitation, you must select either the operating system repository or an LDAP repository for user authentication and authorization in IBM Spectrum Control. If you upgraded from an earlier version of IBM Spectrum Control with both the operating system repository and an LDAP repository configured, you can keep using both repositories. However, it is recommended that you select either the operating system repository or the LDAP repository.

The use of the IBM Spectrum Control single sign-on feature is limited when you change from LDAP authentication to operating system authentication in IBM Spectrum Control. Storage system element managers do not support the operating system repository for single sign-on, even if the element manager is installed on the same system as IBM Spectrum Control.

Procedure

1. Back up the ldapregistry.xml file in the *installation_dir/wlp/usr/servers/webServer/registry/* directory.
2. Log on as an administrator to the IBM Spectrum Control GUI.
3. Click Settings > User Management.
4. Click Edit Authentication.
5. On the Authentication Configuration page, select Local OS and click Save.
6. On the OS Settings page, enter your local operating system user name, password, group name and click Save.

Note: The user name must be a member of the group.

If your change from LDAP authentication to Local OS authentication is successful, you are logged out of the IBM Spectrum Control GUI. You can log into the GUI using your local operating system credentials. You cannot log into the GUI using your LDAP credentials.

If your change from LDAP authentication to Local OS authentication is *not* successful, click Discard. After the confirmation, your previous LDAP authentication is restored and you are returned to the User Management page.

Results

You have changed from LDAP authentication to local operating system in IBM Spectrum Control. You can log on to the IBM Spectrum Control GUI with your local operating system credentials.

Configuring user authentication alternatives

If you encounter any issues with the recommended procedures for configuring user authentication, see the following topics for support.

- [Enabling secure communication between IBM Spectrum Control and the LDAP repository](#)
You can use the Secure Socket Layer (SSL) protocol to secure the communication between IBM Spectrum Control and the LDAP repository that you are using for user authentication. The SSL protocol provides security and data integrity for communications over Transmission Control Protocol/Internet Protocol (TCP/IP) networks.
- [Disabling secure communication between IBM Spectrum Control and the LDAP repository](#)
You can disable the Secure Socket Layer (SSL) protocol between IBM Spectrum Control and the LDAP repository at any time using the IBM Spectrum Control GUI.
- [Exporting SSL certificate from the IBM Security Directory Server to a file](#)
To secure communications between IBM Spectrum Control and IBM® Security Directory Server, you must export the SSL certificate to a file. The file that is created can then be added to the keystore for IBM Spectrum Control.
- [Exporting SSL certificate from the Microsoft Active Directory to a file](#)
To secure communications between IBM Spectrum Control and Microsoft Active Directory, you must export the SSL certificate from the Microsoft Active Directory to a file. The file that is created can then be added to the web server keystore for IBM Spectrum Control.
- [Adding the SSL certificate for the LDAP server to the web server keystore that uses the IKEYCMD command](#)
To secure communications between the IBM Spectrum Control server and the LDAP server, you must add the SSL certificate from the LDAP server to the web server keystore for IBM Spectrum Control.

Enabling secure communication between IBM Spectrum Control and the LDAP repository

You can use the Secure Socket Layer (SSL) protocol to secure the communication between IBM Spectrum® Control and the LDAP repository that you are using for user authentication. The SSL protocol provides security and data integrity for communications over Transmission Control Protocol/Internet Protocol (TCP/IP) networks.

Before you begin

You added an LDAP repository to the federated repositories for IBM Spectrum Control and your system is operating properly with non-secure communication between IBM Spectrum Control and the LDAP repository. Before you implement the following procedure, add the SSL certificate from the LDAP server to the IBM Spectrum Control web server keystore or the connectivity between IBM Spectrum Control and the LDAP server fails.

Procedure

1. Log in to the IBM Spectrum Control GUI as an LDAP user with the Administrator role.
2. In the menu bar, go to Settings > User Management.
3. On the User Management page, click Edit Authentication.
4. On the Authentication Configuration page, click Advanced Configuration Options.
Depending on the LDAP user account that you used to log into the IBM Spectrum Control GUI, you might have to explicitly log into the Liberty Admin Center as the Common User or the file-based user.
5. On the Server Config page, click LDAP User Registry.
6. Change the value in the Ldap ssl enabled field from false to true.
7. Change the value of the Port field to the LDAP server port that listens for secure communications.
The typical value is 636. Depending on your LDAP server configuration, you can specify a different port. If you do not know which port to use, contact your LDAP server administrator.
8. Click Save and log out of the Liberty Admin Center.
9. On the Authentication Configuration page, click Cancel.

Results

Secure communications are established between IBM Spectrum Control and the LDAP repository with SSL protocol.

Related tasks

- [Exporting SSL certificate from the IBM Security Directory Server to a file](#)
- [Exporting SSL certificate from the Microsoft Active Directory to a file](#)
- [Adding the SSL certificate for the LDAP server to the web server keystore that uses the IKEYCMD command](#)

Disabling secure communication between IBM Spectrum Control and the LDAP repository

You can disable the Secure Socket Layer (SSL) protocol between IBM Spectrum® Control and the LDAP repository at any time using the IBM Spectrum Control GUI.

Procedure

1. Log in to the IBM Spectrum Control GUI as an LDAP user with Administrator role.
2. In the menu bar, go to Settings > User Management.
3. On the User Management page, click Authentication Configuration.
4. On the Authentication Configuration page, click Advanced Configuration Options.
Depending on the LDAP user account that you used to log into the IBM Spectrum Control GUI, you might have to explicitly log into the Liberty Admin Center as the Common User or the file-based user.
5. On the Server Config page, click LDAP User Registry.
6. Change the value of the Ldap ssl enabled field from true to false.
7. Change the value of the Port field to the LDAP server port that listens for non-secure communications.
The typical value is 389. Depending on your LDAP server configuration, you can specify a different port. If you do not know which port to use, contact your LDAP server administrator.
8. Click Save and log out of the Liberty Admin Center.
9. On the Authentication Configuration page, click Cancel.

Exporting SSL certificate from the IBM Security Directory Server to a file

To secure communications between IBM Spectrum Control and IBM® Security Directory Server, you must export the SSL certificate to a file. The file that is created can then be added to the keystore for IBM Spectrum Control.

Before you begin

Important: This topic is an example of exporting the SSL certificate from the IBM Security Directory Server to a file.
If your LDAP server is the IBM Security Directory Server verify that the Web Administration tool is installed with your IBM Security Directory Server because it includes the correct IBM Key Management (iKeyman) utility.
For more information about exporting the SSL certificate from the LDAP server, see your LDAP administrator and the documentation for your specific LDAP server product.

Procedure

1. Open the IBM Key Management utility in your IBM WebSphere® Application Server directory structure.
2. Select Key Database File, Open.
3. Complete the following steps:
 - a. In Key database type field, select CMS.
 - b. In the File Name field, click Browse and go to the location of the key database (.kdb) file that is associated with your IBM Security Directory Server.
 - c. Click Open.
 - d. Click OK.
4. On the Password Prompt page, enter the correct password for the key database file and click OK.
5. In the Personal Certificates list, select the certificate that is the SSL certificate for the IBM Security Directory Server and click Extract Certificate.
6. Select Base64-encoded ASCII data as the data type and provide a Certificate file name, Location, and click OK.

Results

The SSL certificate is exported from IBM Security Directory Server to a file so it can be added to the web server keystore for IBM Spectrum® Control.

Exporting SSL certificate from the Microsoft Active Directory to a file

To secure communications between IBM Spectrum Control and Microsoft Active Directory, you must export the SSL certificate from the Microsoft Active Directory to a file. The file that is created can then be added to the web server keystore for IBM Spectrum Control.

Before you begin

Important: This topic is an example of exporting the SSL certificate from the Microsoft Active Directory to a file.

You can use the Certification Authority tool to export the SSL certificate.

For more information about exporting the SSL certificate from the LDAP server, see your LDAP administrator and the documentation for your specific LDAP server product.

Procedure

1. Open the Certification Authority tool.
2. Select Certification Authority, Issued Certificates.
3. Select your current SSL certificate for the Microsoft Active Directory and open the certificate.
4. On the Certificate page, click the Details tab and click Copy to File.
5. In the Certificate Export Wizard, select Base-64 encoded X.509 (.CER) and click Next.
6. Provide a file name and click Next.
7. Review your settings to verify that you have the correct Base-64 encoded X.509 (.CER) file and click Finish.

Results

The SSL certificate is exported from Microsoft Active Directory to a file so it can be added to the web server keystore for IBM Spectrum® Control.

Adding the SSL certificate for the LDAP server to the web server keystore that uses the IKEYCMD command

To secure communications between the IBM Spectrum® Control server and the LDAP server, you must add the SSL certificate from the LDAP server to the web server keystore for IBM Spectrum Control.

About this task

Tip: Contact your LDAP administrator and obtain the SSL certificate for your LDAP server. The SSL certificate must be in the form of a Base64-encoded file.

You can use the IBM® Key Management (IKEYCMD) command to add the LDAP SSL certificate to the IBM Spectrum Control web server keystore.

Procedure

1. Copy the Base64-encoded file to the location of your IBM Spectrum Control server.
2. Log on to the IBM Spectrum Control server with administrative privileges.
3. Open a command prompt and go to *installation_dir\re/bin* directory.
4. Choose one of these options to add the SSL certificate for the LDAP server to the keystore for the web server:
 - Windows operating systems:

```
keycmd -cert -add -db installation_dir\wlp\usr\servers\webServer\resources
\security\key.p12 -pw password -label label -file LDAP SSL certificate
```

- AIX® and Linux® operating systems:

```
./ikeycmd -cert -add -db installation_dir\wlp\usr\servers\webServer/resources
/security/key.p12 -pw password -label label -file LDAP SSL certificate
```

Where the *label* value is for the LDAP SSL certificate you are adding to the IBM Spectrum Control web server keystore. The *password* value is the password that is associated with the keystore. The *default* value for this password is *default*. The *LDAP SSL certificate* value is the Base64-encoded file that contains the SSL certificate from your LDAP server.

5. Restart the IBM Spectrum Control web server.
6. Choose one of these options to verify that the SSL certificate for LDAP was added to the keystore for the web server:
 - For Windows operating systems:

```
ikeycmd -cert -list -db installation_dir\wlp\usr\servers\webServer
\resources\security\key.pl2 -pw password
```

- For AIX and Linux operating systems:

```
./ikeycmd -cert -list -db installation_dir/wlp/usr/servers/webServer
/resources/security/key.pl2 -pw password
```

Results

The SSL certificate from the LDAP server was added to the IBM Spectrum Control web server keystore to enable secure communications.

Related tasks

- [Starting and stopping the IBM Spectrum Control servers](#)
- [Exporting SSL certificate from the IBM Security Directory Server to a file](#)
- [Exporting SSL certificate from the Microsoft Active Directory to a file](#)

Using the ldapEntityType element for advanced LDAP configuration

To narrow the IBM Spectrum® Control view of your LDAP structure so that you can find and map your groups to IBM Spectrum Control roles, you need to configure the **ldapEntityType** element within the **ldapRegistry** element.

Example: LDAP user and group scenario

In this scenario, you want to provide access to LDAP users that reside in a different node of the LDAP structure than your LDAP groups. You also want to prevent authorization of LDAP users and groups that are not associated with IBM Spectrum Control.

In this scenario, the LDAP users use the following distinguished names:

- LDAP user 1: **cn=LDAPUser1,ou=MarketingUsers,dc=storage,dc=company,dc=com**
- LDAP user 2: **cn=LDAPUser2,ou=SalesUsers,dc=storage,dc=company,dc=com**
- LDAP user 3: **cn=LDAPUser3,ou=ManagementUsers,dc=storage,dc=company,dc=com**

In this scenario, the LDAP groups use the following distinguished names:

- LDAP user 1 is a member of LDAP group 1: **cn=LDAPGroup1,ou=MarketingGroups,dc=storage,dc=company,dc=com**
- LDAP user 2 is a member LDAP group 2: **cn=LDAPGroup2,ou=SalesGroups,dc=storage,dc=company,dc=com**
- LDAP user 3 is a member of LDAP group 3: **cn=LDAPGroup3,ou=ManagementGroups,dc=storage,dc=company,dc=com**

In this scenario, **LDAPUser1** and **LDAPUser2** are in different nodes of the LDAP structure than the associated LDAP groups, **LDAPGroup1** and **LDAPGroup2**. You need to find and map **LDAPGroup1** and **LDAPGroup2** to the IBM Spectrum Control roles and be able to log in to IBM Spectrum Control as the **LDAPUser1** and **LDAPUser2**.

When you configure IBM Spectrum Control for LDAP authentication and you encounter this scenario, set the baseDN value in your LDAP XML template file to **dc=storage, dc=company, dc=com**.

In following example, this baseDN value is common to all LDAP users and groups.

Example of the LDAP XML template file that implements this baseDN value:

```
<server description="IBM Web Server">
  <ldapRegistry activatedFilters="active_dir_server" baseDN="dc=storage,dc=company,dc=com"
    bindDN="cn=Administrator,cn=users,dc=storage,dc=company,dc=com" bindPassword="password"
    host="ldap.storage.company.com"
  id="_home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_configuring_cfg_ldap_advanced_LDAP1" ignoreCase="true"
    ldapType="Microsoft Active Directory"
    port="389" realm="TPCRealm" sslEnabled="false">
    </ldapRegistry>

    <activatedLdapFilterProperties groupFilter="(& (cn=%v) (objectcategory=group))" groupIdMap="*:cn"
      groupMemberIdMap="memberof:member"
    id="_home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_configuring_cfg_ldap_advanced_active_dir_server"
      userFilter="(& (sAMAccountName=%v) (objectcategory=user))" userIdMap="user:sAMAccountName"/>

    <federatedRepository>
      <primaryRealm allowOpIfRepoDown="true" name="TPCRealm">
        <participatingBaseEntry name="dc=storage,dc=company,dc=com"/>
        <!-- The next two entries must NOT be changed -->
        <participatingBaseEntry name="o=TPCRealm"/>
        <participatingBaseEntry name="o=OSRealm"/>
      </primaryRealm>
    </federatedRepository>
  </server>
```

When you use this example LDAP XML template file, you also are able to find and map **LDAPGroup3** to an IBM Spectrum Control role. This behavior enables **LDAPUser3** to log in to IBM Spectrum Control, which is not part of the successful outcome. Use the **ldapEntityType** element within the **ldapRegistry** element so you can only find and map **LDAPGroup1** and **LDAPGroup2** to IBM Spectrum Control roles.

Tip: Whenever you modify the **ldapRegistry** element, in the LDAP XML template file, verify that the XML file is valid by opening it in a web browser and checking the results.

This is the previous example of the LDAP XML template file that was edited to use the `ldapEntityType` element within the `ldapRegistry` element:

```
<server description="IBM Web Server">
  <ldapRegistry activatedFilters="active_dir_server" baseDN="dc=storage,dc=company,dc=com"
bindDN="cn=Administrator,cn=users,dc=storage,dc=company,dc=com" bindPassword="password"
host="ldap.storage.company.com"
id="_home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_configuring_cfg_ldap_advanced_LDAP1" ignoreCase="true"
ldapType="Microsoft Active Directory"
port="389" realm="TPCRealm" sslEnabled="false">

    <ldapEntityType name="Group">
      <objectClass>group</objectClass>
      <searchBase>ou=MarketingGroups,dc=storage,dc=company,dc=com</searchBase>
      <searchBase>ou=SalesGroups,dc=storage,dc=company,dc=com</searchBase>
    </ldapEntityType>

    <ldapEntityType name="PersonAccount">
      <objectClass>user</objectClass>
      <searchBase>ou=MarketingUsers,dc=storage,dc=company,dc=com</searchBase>
      <searchBase>ou=SalesUsers,dc=storage,dc=company,dc=com</searchBase>
    </ldapEntityType>

  </ldapRegistry>

  <activatedLdapFilterProperties groupFilter="(& (cn=%v) (objectcategory=group)) "
groupIdMap="*:cn" groupMemberIdMap="memberof:member"
id="_home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_configuring_cfg_ldap_advanced_active_dir_server"
userFilter="(& (sAMAccountName=%v) (objectcategory=user)) "
userIdMap="user:sAMAccountName"/>

  <federatedRepository>
    <primaryRealm allowOpIfRepoDown="true" name="TPCRealm">
      <participatingBaseEntry name="dc=storage,dc=company,dc=com"/>
      <!-- The next two entries must NOT be changed -->
      <participatingBaseEntry name="o=TPCRealm"/>
      <participatingBaseEntry name="o=OSRealm"/>
    </primaryRealm>
  </federatedRepository>
</server>
```

When you use the LDAP XML template that implements the `ldapEntityType` element, it prevents you from finding and mapping `LDAPGroup3` to IBM Spectrum Control role. However, you can find and map `LDAPGroup1` and `LDAPGroup2` to IBM Spectrum Control roles and `LDAPUser1` and `LDAPUser2` can log in to IBM Spectrum Control.

Adding customized text to the logon page

On the logon page for IBM Spectrum® Control, you can show customized text when users access the GUI.

Procedure

1. Open the directory that was created to install IBM Spectrum Control:
 - The default installation directory for Windows operating systems is C:\Program Files\IBM\TPC.
 - The default installation directory for AIX® or Linux® operating systems is /opt/IBM/TPC.
2. Go to one of the following directories:

Windows operating systems
 \wlp\usr\servers\webServer\customization
AIX or Linux operating systems
 /wlp/usr/servers/webServer/customization
3. Open the LoginText.html file in a text editor:
 - a. Type the text that you want to show to the user before they log on to the GUI.
 Tip: To format the text that you want to add, you can use HTML tags, such as paragraph tags, list tags, bold tags, and italic tags.
 - b. Save the LoginText.html file.
4. Open the GUI.

Results

The customized text that you added is shown under the logon page.

Managing a SAN without agents

You can manage a SAN when there are no agents.

About this task

In the following situations, there might not be any agents on the SAN:

- The hosts do not currently have a Storage Resource agent or Fabric agent installed.
- The host operating system is not supported by the Storage Resource agent or Fabric agent.
- The customer requirements do not require the deployment of a Storage Resource agent or Fabric agent.

In these cases, it is recommended that an agent is installed on the Device server itself. This action allows the Device server to use advanced features like Remote Node Identification, which requires an agent.

Normally the Device server does not have a Fibre Channel host bus adapter. In this configuration, the following steps are taken:

1. A Fibre Channel host bus adapter is added to the manager.
2. An agent is installed on the Device server (the Device server is installed first).
3. All storage devices are verified to ensure that they use LUN masking techniques. The LUN masking techniques prevent the Device server from accessing the disks used by the host systems.
4. The Fibre Channel host bus adapter is attached to the SAN to be managed. This host is added to each zone that is intended to be managed by the Device server.

Setting timeout values for the Device server

If a probe or discovery of a storage subsystem times out before the operation completes, you can increase the timeout values for the Device server.

Before you begin

If a probe or discovery of a storage subsystem times out before the operation completes, you receive the following error message:

**HWN021650E Encountered timeout while connecting to CIMOM IP:port.
Check the CIMOM or increase timeout value.**

where *IP* is the IP address, and *port* is the port number. If you determine that the Common Information Model Object Manager (CIMOM) is not the cause of the problem, you can use the command-line interface (CLI) to increase the timeout values for the Device server.

For those storage systems that use native interfaces to connect to IBM Spectrum® Control you see this error message:

HWN020103E The external process exceeded the timeout limit and was cancelled.

The following storage systems use native interfaces to connect to IBM Spectrum Control:

- System Storage® DS8000®
- SAN Volume Controller
- The XIV®
- IBM Spectrum Accelerate
- Storwize® V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified
- IBM FlashSystem® devices that run IBM Spectrum Virtualize
- IBM Spectrum Scale
- IBM® Cloud Object Storage

Procedure

1. Run the **getdscfg** command to determine the current values of the timeout properties. From the command prompt, enter the following command:

```
cli>tpctool getdscfg -user user -pwd password -url host:port  
-property timeout_property
```

where:

- *user* is an IBM Spectrum Control user ID.
- *password* is the password for the IBM Spectrum Control user ID.
- *host* is the host name or IP address, and *port* is a valid port number for the HTTP service of the Device server. The default value for *port* is typically 9550.
- *timeout_property* is one of the following strings:
 - httpTimeout
 - CIMClientWrapper.Timeout
 - Probe.Timeout.Array
 - Probe.Timeout.LMM
 - Discovery.Timeout
 - CIMOMManager.TestConnectionTimeout

Important: Timeout properties are displayed in milliseconds. If the value is 0 (zero), it means that there is no timeout.

For the storage systems that use the native interface, the *timeout_property* strings are:

- NAPI.Timeout.TestConnection
- NAPI.Timeout.Probe
- NAPI.Timeout.EventPoll

2. Run the **setdscfg** command to increase the timeout value. Run the following command:

```
cli>tpctool setdscfg -user user -pwd password -url host:port  
-property timeout_property timeout_value
```

Tip: For Storwize V7000 Unified, the refresh of configuration data from the storage system can take some time to complete and might cause the probe to time out, even if the timeout value is increased. To reduce the duration of the probe run for Storwize V7000 Unified, run the following command:

```
cli>tpctool setdscfg -user user -pwd password -url host:port  
-property Probe.GetRecentStorwizeUnifiedData -context DeviceServer false
```

This command changes the probe configuration to use cached configuration data from the storage system, which reduces the duration of the probe run. However, the information that is collected by the probe might be slightly out of date.

What to do next

For more information about **tpctool**, see [tpctool command](#). You also can view help from the command line by issuing the command with the **-help** option.

Configuring Service Location Protocol

You can enable IBM Spectrum® Control to discover a larger set of storage devices through Service Location Protocol (SLP). In addition to some of the more common SLP configuration issues, there is also information about router configuration, SLP directory agent configuration, and environment configuration.

About this task

For additional information about SLP, see the **Service Location Protocol Request for Comments** website at <http://www.ietf.org/rfc/rfc2165.txt>.

Note: The storage systems that use the native interfaces (DS8000®, XIV®, SAN Volume Controller, and Storwize® V7000) do not use SLP discovery.

- **Router configuration**
Configure the routers in the network to enable general multicasting or to allow multicasting for the SLP multicast address and port, 239.255.255.253, port 427. The routers of interest are the ones associated with subnets that contain one or more storage devices that are to be discovered and managed by IBM Spectrum Control.
- **SLP directory agent configuration**
Review these suggestions when you configure the SLP directory agent.
- **Environment configuration**
This section provides information about the configuration of your environment.
- **SLP registration and the slptool command**
- **SLP discovery**
A common problem with SLP discovery is due to IP multicasting being disabled on the network router. Communication between the SLP SA and UA is done with IP multicasting. Follow these recovery procedures when there are SLP discovery problems and IP multicasting is disabled on the network router.

Router configuration

Configure the routers in the network to enable general multicasting or to allow multicasting for the SLP multicast address and port, 239.255.255.253, port 427. The routers of interest are the ones associated with subnets that contain one or more storage devices that are to be discovered and managed by IBM Spectrum® Control.

About this task

To configure your router hardware and software, refer to your router and configuration documentation.

SLP directory agent configuration

Review these suggestions when you configure the SLP directory agent.

About this task

Configure the SLP directory agents (DAs) to circumvent the multicast limitations. With statically configured DAs, all service requests are unicast by the user agent. Therefore, it is possible to configure one DA for each subnet that contains storage devices that are to be discovered by IBM Spectrum® Control. One DA is sufficient for each of the subnets. Each of these DAs can discover all services within its own subnet, but no other services outside its own subnet. To allow IBM Spectrum Control to discover all the devices, it needs to be statically configured with the addresses of each of these DAs. This operation can be accomplished by using the IBM Spectrum Control Discovery Preference panel.

You can use this panel to enter a list of DA addresses. IBM Spectrum Control sends unicast service requests to each of these statically configured DAs, and sends multicast service requests on the local subnet on which IBM Spectrum Control is installed. Configure an SLP DA by changing the configuration of the SLP service agent (SA) that is included as part of an existing CIM Agent installation. This action causes the program that normally runs as an SLP SA to run as an SLP DA.

Note: The change from SA to DA does not affect the CIMOM service of the subject CIM Agent, which continues to function normally, sending registration and deregistration commands to the DA directly.

Environment configuration

This section provides information about the configuration of your environment.

About this task

It might be advantageous to configure SLP DAs in the following environments:

- In environments where there are other non-IBM Spectrum® Control SLP UAs that frequently perform discovery on the available services, an SLP DA must be configured. This action ensures that the existing SAs are not overwhelmed by too many service requests.
- In environments where there are many SLP SAs, a DA helps decrease network traffic that is generated by the multitude of service replies. It also ensures that all registered services can be discovered by a given UA. The configuration of an SLP DA is recommended when there are more than 60 SAs that need to respond to any given multicast service request.

SLP registration and the slptool command

About this task

IBM Spectrum® Control uses Service Location Protocol (SLP) discovery, which requires that all the discovered CIM agents are registered by using the SLP.

In a non-multicast network environment, SLP can only discover CIM agents that are registered in its IP subnet. For CIM agents outside of the IP subnet, you need to use an SLP DA and register the CIM agents by using **slptool**. Ensure that the CIM_InteropSchemaNamespace and Namespace attributes are specified.

For example, enter the following command:

```
slptool register service:wbem:https://myhost.com:port
```

Where *myhost.com* is the name of the server that is hosting the CIM agent, and *port* is the port number of the service, for example 5989.

Note: **slptool** is installed with a CIM agent. Run the command from the computer that is hosting the CIM agent.

SLP discovery

A common problem with SLP discovery is due to IP multicasting being disabled on the network router. Communication between the SLP SA and UA is done with IP multicasting. Follow these recovery procedures when there are SLP discovery problems and IP multicasting is disabled on the network router.

About this task

Note: The storage systems that use native interfaces, for example, DS8000®, XIV®, SAN Volume Controller, and Storwize® V7000 do not use SLP discovery. There are two recovery procedures when there are SLP discovery problems and IP multicasting is disabled on the network router:

1. Configure one DA for each subnet within the environment.
2. Enable IP multicasting on the router which is disabled by default. Here is a list of common router configurations for multicasting:
 - Internet Group Management Protocol (IGMP) is used to register individual hosts in particular multicast groups and to query group membership on particular subnets.
 - Distance Vector Multicast Routing Protocol (DVMRP) is a set of routing algorithms that use a technique called reverse path forwarding. These algorithms provide the best solution for how multicast packets are to be routed in the network.
 - Protocol-Independent Multicast (PIM) comes in two varieties: dense mode (PIM-DM) and sparse mode (PIM-SM). The dense mode and sparse mode routines are optimized for networks where either a large percentage of nodes requires multicast traffic (dense) or a small percentage of nodes requires the sparse traffic.
 - Multicast Open Shortest Path First (MOSPF) is an extension of OSPF. It is a link-state unicast routing protocol that attempts to find the shortest path between any two networks or subnets to provide the most optimal packet routing.

To properly configure the routers for multicasting, see the reference and configuration documentation from the router manufacturer.

Configuring IP addressing

This section provides information about configuring IP addressing.

About this task

- **Configuring IBM Spectrum Control with multiple IP addresses**
If the system where IBM Spectrum Control is to be installed has multiple IP addresses, then a configuration value must be set manually as a post-installation task by using the **tpctool setdscfg** command. The value to be set is for the local IP address, which must be used for subscription for CIM Indications for CIM agents.
- **Changing the HOSTS file**
When you install IBM Spectrum Control on your Windows operating systems, you must follow these steps to avoid addressing problems with the systems you want to manage. These problems are caused by the address resolution protocol that returns the host short name rather than the fully qualified host name. You can avoid this by modifying the entries in the corresponding host tables on the DNS server and on the local computer system. The fully qualified host name must be listed before the short name in each entry that is associated with systems managed by IBM Spectrum Control.

Configuring IBM Spectrum Control with multiple IP addresses

If the system where IBM Spectrum® Control is to be installed has multiple IP addresses, then a configuration value must be set manually as a post-installation task by using the **tpctool setdscfg** command. The value to be set is for the local IP address, which must be used for subscription for CIM Indications for CIM agents.

About this task

Restriction: This task does not apply to storage systems that use the native interfaces, for example, DS8000®, XIV®, SAN Volume Controller, and Storwize® V7000. If you are using IPv6 computers, see [Planning for Internet Protocol Version 6](#).

For multiple IPv6 addresses, the IPv6 address to use for CIM indication subscription by IBM Spectrum Control can be specified by setting the property **System.LocalIPv6Address** as described.

With dual stack IPv4 and IPv6 IBM Spectrum Control servers, two IP addresses are required to subscribe to IPv4 CIMOMs and IPv6 CIMOMs. The configuration property **System.LocalIPv6Address** is used for IPv6 CIMOMs and the property **System.LocalIPAddress** is used for IPv4 CIMOMs.

To change the IP address, follow these steps:

Procedure

1. Open a command prompt window on the server system.
2. Change to the following directory:

```
cd installation_dir\cli
```

3. Enter the following command:

```
tpctool setdscfg -user user_ID -pwd password -url  
host:port -property System.LocalIPv6Address value
```

Where:

user_ID

Is the user ID.

password

Is the password for the user.

host

Is either the host name or IP address of the system that is running IBM Spectrum Control.

port

Is a valid port number for the HTTP service of the Device server (the default is 9550).

value

Is the local IP address, which must be used for subscription for CIM Indications for CIM agents.

4. Verify that the command was successful by entering this command:

```
tpctool getdscfg -user user_ID -pwd password  
-url host:port -property System.LocalIPv6Address
```

Changing the HOSTS file

When you install IBM Spectrum® Control on your Windows operating systems, you must follow these steps to avoid addressing problems with the systems you want to manage. These problems are caused by the address resolution protocol that returns the host short name rather than the fully qualified host name. You can avoid this by modifying the entries in the corresponding host tables on the DNS server and on the local computer system. The fully qualified host name must be listed before the short name in each entry that is associated with systems managed by IBM Spectrum Control.

About this task

The **HOSTS** file is in the **%SystemRoot%\system32\drivers\etc** directory. To change the HOSTS file, follow these steps:

Procedure

1. Open the HOSTS file in a text editor.
2. Add, remove, or modify the host entries. In the following example of a HOSTS file, the short name is incorrectly listed before the fully qualified host name. This can cause address resolution problems in IBM Spectrum Control.

```
# Copyright (c) 1993-2009 Microsoft Corp.  
#  
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.  
#  
# This file contains the mappings of IP addresses to host names. Each  
# entry should be kept on an individual line. The IP address should  
# be placed in the first column followed by the corresponding host name.  
# The IP address and the host name should be separated by at least one  
# space.  
#  
# Additionally, comments (such as these) may be inserted on individual  
# lines or following the machine name denoted by a '#' symbol.  
#  
# For example:  
#  
#       102.54.94.97       rhino.acme.com   # source server  
#       38.25.63.10       x.acme.com      # x client host  
#  
192.168.123.146          jason           jason.groupa.mycompany.com
```

3. In the following example, the order of the host names has been changed so that the fully qualified host name is placed before the short name. The host names must be entered in the order that is shown so IBM Spectrum Control can locate the host. Use this format for any hosts that are associated with IBM Spectrum Control.

```
# For example:  
#  
#       102.54.94.97       rhino.acme.com   # source server  
#       38.25.63.10       x.acme.com      # x client host  
#  
192.168.123.146          jason.groupa.mycompany.com   jason
```


Note: Host names are case-sensitive. This is a WebSphere® requirement. For example, if your computer shows the name as JASON (uppercase), then you must enter JASON in the HOSTS file.

Deploying Storage Resource agents

You can manage your Storage Resource agent deployments.

Deploy Storage Resource agents through the user interface rather than a separate installation wizard. You can have only one agent per host that points to the same IBM Spectrum® Control server.

Before you begin: Before you deploy Storage Resource agents, see [Deployment guidelines and limitations for Storage Resource agents](#) for a list of considerations. If you deploy Storage Resource agents on multiple computers at the same time, the computers must have the same administrative user ID and password. IBM Spectrum Control uses these user credentials to log on to the computers when the Storage Resource agents are deployed.

To deploy Storage Resource agents, complete the following steps:

1. In the menu bar, go to Servers > Servers.
2. Click Add Server.
3. Select Deploy an agent for full server monitoring.
4. Select one of the following methods for adding a server:
 - Add a server by manually entering information about the server and the Storage Resource agent.
 - Add one or more servers by importing configuration information from a comma-delimited file.
5. Configure deployment information for the Storage Resource agents.
6. Schedule the agent deployment and the data collection for the servers.
7. Click Finish to deploy the Storage Resource agents.

- [Deployment guidelines and limitations for Storage Resource agents](#)

You must consider the following guidelines and limitations when you manage Storage Resource agents in your environment.

- [Creating a certificate for SSH protocol](#)

Before you install the Storage Resource agents by using the SSH protocol, you can optionally create a certificate.

- [Replacing default SSL certificates for the Data server and Storage Resource agents with custom SSL certificates](#)

IBM Spectrum Control provides default SSL certificates for communication between the Data server and Storage Resource agent. You can replace the default SSL certificates. You must use the script that is provided by IBM Spectrum Control to generate new SSL certificates. You cannot use any third-party tools to generate the custom SSL certificates.

Deployment guidelines and limitations for Storage Resource agents

You must consider the following guidelines and limitations when you manage Storage Resource agents in your environment.

Use the following information when you deploy Storage Resource agents:

- [Multiple Storage Resource agents that are probing or scanning the same storage resources](#)
- [Platforms that support the deployment of Storage Resource agents](#)
- [Product functions that are not available for storage devices monitored by Storage Resource agents](#)
- [Required authority for deploying Storage Resource agents](#)
- [Orphan zones](#)
- [Firewalls and Storage Resource agents deployments](#)
- [Deploying Storage Resource agents on multiple computers](#)
- [Communication between the IBM Spectrum® Control server and a Storage Resource agent](#)
- [Daemon and non-daemon services](#)
- [Port numbers for Storage Resource agents deployed as a daemon service](#)
- [Authentication between the IBM Spectrum Control server and a Storage Resource agent](#)
- [Replacing default SSL certificates](#)
- [Storage Resource agents on the same computer](#)
- [Time zones for computers monitored by Storage Resource agents](#)
- [Connections for Linux® and AIX® operating systems by using Remote Shell protocol \(RSH\)](#)
- [Deployments on Windows - NetBIOS setting](#)
- [Deployments on Windows - User Account Control \(UAC\) remote restrictions](#)

Multiple Storage Resource agents that are probing or scanning the same resources

If multiple Storage Resource agents are set up to probe or scan the same storage resources, the Storage Resource agent that was added to IBM Spectrum Control first is used for the probe or scan. Therefore, only data that is gathered by the first Storage Resource agent is shown.

Platforms that support the deployment of Storage Resource agents

For a list of platforms on which you can deploy Storage Resource agents, see the [IBM Spectrum Control interoperability matrix](#) and go to the *Agents, Servers and Browsers* section.

Product functions that are unavailable for resources that are monitored by Storage Resource agents

Before you deploy a Storage Resource agent, ensure that the product functions you want to use on the monitored resources are available for those agents. The following functions are not available for resources that are monitored by Storage Resource agents:

- Certain relational database monitoring. For list of relational databases that can be monitored by Storage Resource agents, see the [IBM Spectrum Control interoperability matrix](#) and go to the *Agents, Servers and Browsers* section.
- The reporting of HBA, fabric topology, or zoning information for fabrics that are connected to hosts that are running Linux on IBM® System z® hardware. These limitations also apply to Storage Resource agents on all guest operating systems for VMware configurations.

Required authorities for deploying and running Storage Resource agents

Before you can create deployment schedules and deploy Storage Resource agents on target computers, you must meet the following requirements:

- To create deployment schedules, you must be logged in to IBM Spectrum Control with a user ID that has the Administrator role. For information about user roles, see [Authorizing users](#).
- To deploy Storage Resource agents on target computers, you must provide a user ID that has administrative rights on those computers. You enter this ID when you create a deployment schedule. IBM Spectrum Control uses this ID to log on to the target computers and install and configure the necessary runtime files for the agents.

The user under which a Storage Resource agent (daemon or non-daemon) runs must have the following authorities on the target computers:

- On the Linux or AIX operating systems, the user must have root authority. By default, an agent runs under the user 'root'.
- On the Windows operating systems, the user must have Administrator authority and be a member of the Administrators group. By default, a Storage Resource agent runs under the 'Local System' account.

Orphan zones

Storage Resource agents do not collect information about orphan zones. An orphan zone is a zone that does not belong to at least one zoneset.

Firewalls and Storage Resource agent deployments

Before you can deploy a Storage Resource agent on a computer, you must turn off the firewall on that computer. If you do not turn off the firewall, the deployment fails.

Deploying Storage Resource agents on multiple computers

If you deploy Storage Resource agents on multiple computers at the same time, the computers must have the same administrative user ID and password. IBM Spectrum Control uses these user credentials to log on to the computers when you install Storage Resource agents.

Tip: When you deploy Storage Resource agents on multiple computers, a globally unique identifier (GUID) is created for each computer (if one does not exist).

Communication between the IBM Spectrum Control server and a Storage Resource agent

The IBM Spectrum Control server connects to a monitored computer when a Storage Resource agent is deployed and whenever a data collection schedule runs against that agent.

During deployment, the server communicates with the target computer by using one of the following protocols:

- Windows server message block protocol (SMB)
- Secure Shell protocol (SSH)
- Remote execution protocol (REXEC)
- Remote shell protocol (RSH)

After deployment, the type of communication between the server and agent on that computer depends on whether you deployed the agent as daemon service or non-daemon service.

Daemon and non-daemon services

You can deploy a Storage Resource agent as a daemon or non-daemon service:

- A Storage Resource agent that is deployed as a daemon service runs in the background on the monitored computer and listens for requests from the IBM Spectrum Control server. Connectivity between the server and agent is established by using SSL. The server and agent have their respective certificates and no additional information is required besides those certificates and the security that is provided by the SSL protocol.
- A Storage Resource agent deployed as a service on demand (non-daemon service) runs as a stand-alone executable file on the monitored computer. Communication from the server to the agent uses the same protocol that was used during the deployment of the agent. Communication from the agent to the server uses SSL.
- A Storage Resource agent that is deployed as a daemon service on AIX, Linux, and Windows servers monitors disk paths in near real-time to detect errors. When deployed as a daemon service on an AIX server, the agent also monitors disk error events in near real-time. If the Storage Resource agent detects path status changes or disk errors, they are included in the status of the disks and paths. You can define alerts so that you are notified of changes to the status of the paths on monitored disks.

Only status changes for existing paths are detected. If a new path is added, or an existing path is removed, the number of paths that is displayed is not updated immediately. The number of paths is updated after the next scheduled probe collects data.

If a disk on an AIX server has an error status and you fix the error, you might want the new status of the disk to be displayed immediately. To display the new status immediately, you must reset the status indicator for the disk. To reset the status indicator, use the **errclear** command to clear the error log. To clear the error log, use the following syntax:

```
errclear -d H -N disk_name 0
```

For example, if you fixed an error on hdisk4, and want to display the new status immediately, run the following command:

```
errclear -d H -N hdisk4 0
```

If you do not reset the status indicator for the disk, the status changes automatically after a few hours.

For information about the **errclear** command, see [errclear Command](#).

Port numbers for Storage Resource agents deployed as a daemon service

The following port numbers are used by Storage Resource agents that are deployed as daemon service:

- 9567 (For the Storage Resource agent that is deployed on the same server as IBM Spectrum Control.)
- 9510 (For Storage Resource agents that are deployed on remote servers.)

Storage Resource agents that are deployed as a non-daemon service do not use a port.

Authentication between the IBM Spectrum Control server and a Storage Resource agent

IBM Spectrum Control requires the correct authentication information (user name, password, port, certificate location, or passphrase) for monitored computers each time it communicates with Storage Resource agents on those computers. If the authentication information changes for a host computer on which a Storage Resource agent is deployed, the authentication information for that agent must be updated by using the Modify Agents > Update Credentials action on the Servers page in the GUI.

Replacing default SSL certificates

IBM Spectrum Control provides default SSL certificates for communication between the Data server and Storage Resource agent.

IBM Spectrum Control 5.2.2 uses SSL certificates with 2048-bit encryption keys whereas previous versions of IBM Spectrum Control used 1024-bit encryption keys. If you upgrade IBM Spectrum Control from a version earlier than 5.2.2, your SSL certificates are not updated automatically. If you want to use 2048-bit encryption keys with previous versions of IBM Spectrum Control, you must replace the default SSL certificates with custom SSL certificates.

For information about how to replace SSL certificates, see [Replacing default SSL certificates for the Data server and Storage Resource agents with custom SSL certificates](#).

Storage Resource agents on the same computer

You cannot deploy a Storage Resource agent on a computer where a Storage Resource agent is already installed and pointing to the same Data server. You can deploy a Storage Resource agent on the same computer as another Storage Resource agent if those agents communicate with different Data servers and use different ports when you listen for requests.

Time zones for computers that are monitored by Storage Resource agents

The time zones of computers that are monitored by Storage Resource agents are shown as Greenwich mean time (GMT) offsets in IBM Spectrum Control reports. For example, a computer in Los Angeles shows the following time zones in the By Computer report in Asset reporting:

(GMT-8:00) GMT-8:00

Connections for Linux and AIX operating systems by using Remote Shell protocol (RSH)

If RSH is configured to use a user ID and password, the connection fails. To successfully connect to a system by using RSH, you must set up the `.rhosts` file (in the home directory of the account). RSH must be configured to accept a login from the system that is running your application.

Deployments on Windows operating systems - NetBIOS setting

To install a Storage Resource agent on Windows targets, the Enable NetBIOS over TCP/IP option must be selected in the Control Panel settings for the computer's network connections properties. To set this option, complete the following steps:

1. Open Windows Control Panel. For information about how to open Windows Control Panel, see [Accessing administration tools](#).
2. Select Network and Dial-Up Connections, > some_connection > Properties > Internet Protocol (TCP/IP) > Advanced > WINS > Enable NetBIOS over TCP/IP.

To determine whether these ports are not blocked for inbound requests, see the documentation for your firewall.

To determine whether security policies are blocking the connection ports, open Administrative Tools. For information about how to open Administrative Tools, see [Accessing administration tools](#).

Depending on whether your policies are stored locally or in Active Directory, follow these directions:

Policies that are stored locally

For policies that are stored locally, complete the following steps:

1. Open Windows Administrative Services.
2. Click Local Security Policy > IP Security Policies on Local Computer.

Policies that are stored in Active Directory

For policies that are stored in Active Directory, examine the IP security policies and edit or remove filters that block the ports:

- Click Administrative Tools > Default Domain Security Settings > IP Security Policies on Active Directory.
- Click Administrative Tools > Default Domain Controller Security Settings > IP Security Policies on Active Directory.

For all Windows systems, the Server service must be running to connect to a Windows system by using the Windows protocol.

The following table lists the ports that are reserved for NetBIOS. Ensure that these ports are not blocked.

Port	Description
135	NetBIOS Remote procedure call. (Not currently used.)
137	NetBIOS name service.
138	NetBIOS datagram. (Not currently used.)
139	NetBIOS session (for file and print sharing).
445	CIFS (on Windows XP).

For Windows, shares must be shared for the Guest or Everyone accounts, and password protected sharing must be disabled. To disable password protected sharing, follow these steps:

1. Click Control Panel > Networking and Sharing Center.
2. Click Change advanced sharing settings.
3. Click the down arrow next to All Networks.
4. Select Turn off password protected sharing.
5. Click Save Changes.
6. Exit from the Control Panel.

Deployments on Windows - User Account Control (UAC) remote restrictions

To install Storage Resource agents remotely on a Windows operating system, you must disable the User Account Control (UAC) remote restrictions on the Windows operating system. User Account Control is a security component on Windows operating systems.

Tip: To disable UAC restrictions, you must modify the computer registry. Serious problems might occur if you modify the registry incorrectly. Therefore, make sure that you follow these steps carefully. For added protection, back up the registry before you modify it. Then, you can restore the registry if problems occur. For information about how to back up and restore the registry, see <http://support.microsoft.com/kb/322756/>.

To disable UAC remote restrictions, follow these steps:

1. Open the Windows Run window. For information about how to open the Run window, see [Accessing administration tools](#).
2. Enter **regedit** and click OK.
3. Locate and click the following registry subkey:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\ Policies\System

4. Double click the EnableLUA registry entry.
5. In the Edit DWORD (32-Bit) dialog, change the value in the Value data field from 1 to 0.
6. Click OK.
7. Exit the registry editor.

Related reference

- [Planning for Storage Resource agents](#)

Creating a certificate for SSH protocol

Before you install the Storage Resource agents by using the SSH protocol, you can optionally create a certificate.

Note: The Storage Resource agent only supports either DES-EDE3-CBC encryption or no encryption for the private key used in SSH protocol communication between the server and agent. The default encryption that is used in the **ssh-keygen** command on UNIX is always DES-EDE3-CBC. However, with Windows Cygwin, the **ssh-keygen** command generates a key with AES-128-CBC encryption if a passphrase is specified. If there is no passphrase, the private key is generated without encryption. For more information about encryption, see <https://www.openssl.org/docs/man1.0.2/man1/enc.html>.

Creating a certificate for SSH protocol (non-Windows)

The Storage Resource agent only supports either DES-EDE3-CBC encryption or no encryption for the private key used in SSH protocol communication between the server and agent. The default encryption used in the **ssh-keygen** command on UNIX is always DES-EDE3-CBC but with Windows Cygwin, it is using AES-128-CBC encryption if a passphrase is specified. If there is no passphrase, the private key is generated without encryption.

To create a certificate for SSH protocol, complete the following steps:

1. Telnet to the remote machine using the root user ID.
2. To create an SSH certificate on AIX®, you must first install the following packages (if not already installed):

```
openssl.base.openssh.base.client  
openssl.base.server
```

3. Go to the directory where you want to create the certificate:

```
cd ~/.ssh
```

4. Enter **ssh-keygen -t rsa**. Accept the default names (for example, **id_rsa**).
5. Enter the passphrase.
6. Two files are created:

```
id_rsa  
    The private key.  
id_rsa.pub  
    The public key.
```

7. Create an **authorized_key** file in the same location as **id_rsa.pub** by entering the following command:

```
cat id_rsa.pub >> authorized_keys
```

8. Copy the **id_rsa** (private key) to your server machine. For example, to copy the **id_rsa** file to `:\keys\id_rsa` on the IBM Spectrum® Control server (user responses are in boldface type):

```
# ssh-keygen  
Generating public/private rsa key pair.  
Enter file in which to save the key (//.ssh/id_rsa):  
Enter passphrase (empty for no passphrase):  
Enter same passphrase again:  
Your identification has been save in //.ssh/id_rsa.  
Your public key has been save in //.ssh/id_rsa.pub.  
The key fingerprint is:  
xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx root@server  
# cat id_rsa >> authorized_keys  
# ls -l  
total 24  
-rw-r--r-- 1 root system 1743 Oct 15 09:40 authorized_keys  
-rw--- 1 root system 1743 Oct 15 09:39 id_rsa  
-rw-r--r- 1 root system 399 Oct 15 09:39 id_rsa.pub  
#
```

Note: You must copy the file in binary mode.

9. To connect to the remote system by using the private key, enter the following information in the Remote Agent Machines window of the GUI, when you install the Storage Resource agent:
 - User
 - Certificate Location (c:\keys\id_rsa)
 - Passphrase

Setting up an SSH daemon on Windows

On Windows you must run the **ssh-host-config** command.

Note: Cygwin is not a prerequisite for the Storage Resource agent on Windows. To use the SSH protocol on Windows, an SSH software program must be used because Windows does not come with an SSH service. Cygwin is a free software program providing SSH access to a Windows server. Cygwin can be used if you want to run the Storage Resource agent by using the SSH protocol.

You must be in a Cygwin window or be an X term user to create the **sshd** service. In most cases, you click the **cygwin.bat** file to start the Bash shell.

Complete the following steps:

1. Install Cygwin.
2. Set up your **sshd** service in Cygwin.
3. Create the certificate.

Installing Cygwin

To install Cygwin, go to <http://cygwin.com>. This page contains a link that displays help for the setup program and a link to download the setup program. Read the help before running the setup program. Then download the Cygwin program by clicking the **Install Cygwin now** link. Start the setup program on your computer by running the **setup.exe** program. Select the appropriate download option (**Install from Internet**, **Download from Internet**, or **Install from Local Directory**) as described in the help files.

If you are upgrading from an older version of Cygwin to a newer version, you need to remove the **sshd** service before installing the new version of Cygwin.

Accept the default installation options as they are presented to you (Root Directory, Install For, Default Text File Type, and so on). Select a download mirror that is geographically close to your location. Some sites require an FTP account before you can install Cygwin. You can either request an account or simply select another mirror.

During the installation process, a Select Packages list is displayed. Expand the plus sign (+) next to the Admin category and select **cygrunsrv** and the **Bin** check box. Expand the plus sign (+) next to the Net category and select **openssh**. Expand the plus sign (+) next to the Util category and select **diffutils**. Click **Next** to resume the setup program. The time required to download the packages depends on how busy the mirror is, and on the speed of your internet connection. With **openssh** and **cygrunsrv**, the downloaded files require approximately 70 MB of disk space. Allow 20 minutes to 30 minutes for the download and installation to complete.

Setting up your sshd service in Cygwin

Here is an example of the sequence of steps and responses. The responses to the prompts are in boldfaced type.

1. Run the **ssh-host-config** command.

Note: With Cygwin, you might experience permission problems when running the **ssh-host-config** command. If you have permission problems, run these commands:

```
chmod +r /etc/passwd
chmod +r /etc/group
chmod 777 /var
```

```
$ ssh-host-config
```

```
*** Info: Generating missing SSH host keys
*** Query: Overwrite existing /etc/ssh_config file? (yes/no) yes
*** Info: Creating default /etc/ssh_config file
*** Query: Overwrite existing /etc/sshd_config file? (yes/no) yes
*** Info: Creating default /etc/sshd_config file

*** Info: StrictModes is set to 'yes' by default.
*** Info: This is the recommended setting, but it requires that the POSIX
*** Info: permissions of the user's home directory, the user's .ssh
*** Info: directory, and the user's ssh key files are tight so that
*** Info: only the user has write permissions.
*** Info: On the other hand, StrictModes don't work well with default
*** Info: Windows permissions of a home directory mounted with the
*** Info: 'noacl' option, and they don't work at all if the home
*** Info: directory is on a FAT or FAT32 partition.
*** Query: Should StrictModes be used? (yes/no) no
*** Info: Updating /etc/sshd_config file

*** Query: Do you want to install sshd as a service?
*** Query: (Say "no" if it is already installed as a service) (yes/no) yes
*** Query: Enter the value of CYGWIN for the daemon: [] ntsec
*** Info: On Windows Server 2003, Windows Vista, and above, the
*** Info: SYSTEM account cannot setuid to other users -- a capability
*** Info: sshd requires. You need to have or to create a privileged
*** Info: account. This script will help you do so.

*** Info: It's not possible to use the LocalSystem account for services
*** Info: that can change the user id without an explicit password
*** Info: (such as passwordless logins [e.g. public key authentication]
*** Info: via sshd) when having to create the user token from scratch.
*** Info: For more information on this requirement, see
*** Info: https://cygwin.com/cygwin-ug-net/ntsec.html#ntsec-nopasswdl

*** Info: If you want to enable that functionality, it's required to create
*** Info: a new account with special privileges (unless such an account
*** Info: already exists). This account is then used to run these special
*** Info: servers.

*** Info: Note that creating a new user requires that the current account
*** Info: have Administrator privileges itself.

*** Info: No privileged account could be found.

*** Info: This script plans to use 'cyg_server'.
*** Info: 'cyg_server' will only be used by registered services.
*** Query: Do you want to use a different name? (yes/no) no
*** Query: Create new privileged user account 'local_address\cyg_server'
*** Query: (Cygwin name: 'cyg_server')? (yes/no) yes
*** Info: Please enter a password for new user cyg_server. Please be sure
*** Info: that this password matches the password rules given on your system.
*** Info: Entering no password will exit the configuration.
*** Query: Please enter the password:password
*** Query: Reenter:password

*** Info: User 'cyg_server' has been created with password 'password'.
*** Info: If you change the password, please remember also to change the
*** Info: password for the installed services which use (or will soon use)
*** Info: the 'cyg_server' account.

*** Info: The sshd service has been installed under the 'cyg_server'
*** Info: account. To start the service now, call 'net start sshd' or
*** Info: 'cygrunsrv -S sshd'. Otherwise, it will start automatically
*** Info: after the next reboot.
```

```
*** Info: Host configuration finished. Have fun!
```

2. Start the **sshd** service:

- Open a command prompt window.
- Enter **net start sshd** or in a Bash prompt, enter **cygrunsrv -S sshd**.
- Verify that the daemon is running.
- Enter **ps -a**. Examine the output to see if `/usr/sbin/sshd` is contained in the list of running processes.

To stop the service from a Windows command prompt, enter **net stop sshd**. Alternatively, you can change to the `C:\cygwin\bin` directory (or open a Bash shell) and enter **cygrunsrv -E sshd**.

3. When you have started the **sshd** service, test it by entering the following command from a Bash shell prompt:

```
ssh localhost
or
ssh host_name
```

If **localhost** does not work, use the short host name. If you receive a message indicating that the authenticity of **localhost** cannot be established, answer **Yes** to the question "Are you sure you want to continue connecting?" When prompted for your account password on **localhost**, enter the password you use when logging in to the computer.

4. Set the **TEMP** environment variable. For information about setting the environment variable, see <http://www.cygwin.com/cygwin-ug-net/setup-env.html>.

Here is an example of setting the environment variable:

- Click **My Computer** > **Properties** > **Advanced** > **Environment Variables**.
- Under **System variables**, find out the value of **TEMP**. For example, "`C:\WINNT\TEMP`".
- Set the **TEMP** environment variable to point to the Cygwin format of **TEMP** in the `~/ .bashrc` file. For example run the following command:

```
export TEMP=/cygdrive/c/WINNT/temp
```

Uncomment and modify this line in the `~/ .bashrc` file from the default:

```
# export TEMP=/tmp
to
export TEMP=/cygdrive/c/WINNT/temp
```

The Cygwin **sshd** service must be added as a service that starts automatically. To verify this step, click **Start** > **Settings** > **Control Panel** > **Administrative Tools** > **Services**. Look for **CYGBIN sshd** in the name list. Verify that it is started and configured to start automatically.

Creating the certificate

To create a certificate for SSH protocol, complete the following steps:

1. Run this command:

```
cd ~/.ssh
```

2. Generate the public and private keys with a passphrase. The passphrase is required.

From the Bash shell prompt, here is an example of the input and output (user responses are in boldface type):

```
Administrator ~/.ssh
$ openssl genrsa -des3 -out key 1024
Generating RSA private key, 1024 bit long modulus
.....+++++
.....+++++
e is 65537 (0x10001)
Enter pass phrase for key:passphrase
Verifying - Enter pass phrase for key:passphrase

Administrator ~/.ssh
$ chmod 600 ~/.ssh/key
$ ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/home/Administrator/.ssh/id_rsa): key_pairs
Enter passphrase (empty for no passphrase):passphrase
Enter same passphrase again:passphrase
Your identification has been saved in key_pairs.
Your public key has been saved in key_pairs.pub.
The key fingerprint is:
SHA256:ew0Octa24Qw917tRqPcn9hETlRakksKcTgGrPkh4UZs Sheila@IBM243-PC0CJ5EF
The key's randomart image is:
+---[RSA 2048]-----+
|  . . . . .o+ |
|  . o + o . .o |
|  . E . * o . . |
|  . . . oo . . . |
|  . o . . S.B . oo |
|  o o + O B . oo |
|  . o . * o + . |
|  . . . .o+o |
|  . . . .o+ |
+---[SHA256]-----+
Administrator ~/.ssh
$ cat id_rsa.pub >> authorized_keys
$
```

3. Copy the **id_rsa** (private key) to the IBM Spectrum Control server.

4. To connect to the remote system by using the private key, enter the following information in the GUI, when you install the Storage Resource agent:

- User
- Certificate Location (`c:\keys\id_rsa`)

- Passphrase

Replacing default SSL certificates for the Data server and Storage Resource agents with custom SSL certificates

IBM Spectrum® Control provides default SSL certificates for communication between the Data server and Storage Resource agent. You can replace the default SSL certificates. You must use the script that is provided by IBM Spectrum Control to generate new SSL certificates. You cannot use any third-party tools to generate the custom SSL certificates.

Overview of replacing default SSL certificates for the Data server and Storage Resource agents

IBM Spectrum Control uses SSL certificates for communication between the Data server and Storage Resource agents. IBM Spectrum Control provides default SSL certificates for this communication. If you want to generate new certificates, you can replace the default SSL certificates with updated SSL certificates.

Data server certificate

The IBM Spectrum Control Data server uses the **TPCDataServer.jks** and **server.pwd** files for communication with the Storage Resource agents. If you use custom SSL certificates, you must replace these files.

Storage Resource agent certificate

The Storage Resource agent uses the **sra.pem** and **sra.pwd** files for communication with the Data server. These two files are compressed into the **certs.zip** file on the IBM Spectrum Control server system for Storage Resource agent deployment purposes. If you use custom SSL certificates, you must replace these files.

These general steps are for replacing the default SSL certificates:

1. Generate the custom SSL certificates.
2. Stop the Data server and all Storage Resource agents, including the one on the IBM Spectrum Control server.
3. Replace the default SSL certificate for the Data server and all Storage Resource agents. Also, replace the default SSL certificate for the Storage Resource agents in the IBM Spectrum Control installation image or in the Storage Resource agent installation image.
4. Start the Data server and all Storage Resource agents, including the one on the IBM Spectrum Control server.

Important: When you generate custom SSL certificates, the certificates have a start date, end date, and time when they are valid. These dates and times are related to the system where these custom certificates were generated (which is usually the server system). When you install a Storage Resource agent on a remote system, you must check the date and time on the Storage Resource agent system. If the server and agent systems are in the same time zone, they must have the same date and time. Otherwise, the time zone difference must be set.

For example, if the server system is 8:00 PM, the agent system must also be 8:00 PM. If the agent system is set at a different time (for example, 6:00 PM) at the time the SSL custom certificates are generated on the server system with a time of 8:00 PM, the deployment of the Storage Resource agent fails.

How to generate custom SSL certificates

The **createSRACerts.sh** script (for Linux® or UNIX) or the **createSRACerts.bat** file (for Windows) is located in the following directory:

installation_dir/data/sra/tools/certs


Where **installation_dir** is the directory where the IBM Spectrum Control servers are installed. The default directory is **/opt/IBM/TPC** for Linux or UNIX or **C:\Program Files\IBM\TPC** for Windows.

To replace the default SSL certificates, follow these steps:

1. Create the custom SSL certificates.

The **createSRACerts** script creates the custom SSL certificates.

The syntax is:



output_directory

Directory where the certificates are created. You must provide a valid directory. The script creates the **sra_certs_out** subdirectory and places the certificate files in that subdirectory.

rootCAPassword

Root CA password (root certificate authority password). You can enter a new root certificate authority password or you can enter the default root certificate authority password: **s5umEvApR6cafruhustu**.

server_key_password

Server key password. You can enter a new server key password or you can enter the default server key password: **drUtaxahaswefraf9uth**.

server_store_password

Server store password. You can enter a new server store password or you can enter the default server store password: **wr4d5Xekaqafeh5u2a**.

agent_password

Agent password. You can enter a new agent password or you can enter the default agent password: **jawUchezuthew6azEjef**.

Important: The **createSRACerts** script strictly assumes the order of the command line parameters **output_directory**, **rootCAPassword**, **server_key_password**, **server_store_password**, and **agent_password**. For example, if you want to pass the **rootCAPassword** parameter to the script, the **rootCAPassword** parameter must be the second argument to the script and you must also pass the **output_directory** parameter as the first argument to the script.

Another example: If you want to pass the **server_store_password** parameter to the script, the **server_store_password** parameter must be the fourth argument to the script and you must also pass the **server_key_password** parameter as the third argument, the **rootCAPassword** parameter as the second argument, and the **output_directory** parameter as the first argument to the script.

Important: During the script generation, the script prompts you twice for the pass phrase for `tpcrootca.key`. If you enter a new root certificate authority password on the command line when you run the script, enter that same new root certificate authority password at each prompt. If you enter the default root certificate authority password on the command line when you run the script or you do not enter the root certificate authority password on the command line at all when you run the script, enter the default root certificate authority password at each prompt. The following example creates the SSL certificates by using the default passwords and placing the certificate files in the `sra_certs_out` subdirectory of the current working directory:

```
createSRACerts .
```

The following examples create the SSL certificates by using the default passwords and placing the certificate files in `C:\Temp\sra_certs_out\` on Windows and in `/tmp/sra_certs_out/` on UNIX or Linux.

Windows

```
createSRACerts C:\temp
```

UNIX or Linux

```
./createSRACerts.sh /tmp
```

The following examples create the SSL certificates by using new passwords for the root certificate authority password and the server key password and placing the certificate files in the `C:\Temp\sra_certs_out\` directory on Windows and in the `/tmp/sra_certs_out/` directory on UNIX or Linux:

Windows

```
createSRACerts C:\temp newpasswordforrootCA newpasswordforserver
```

UNIX or Linux

```
./createSRACerts.sh /tmp newpasswordforrootCA newpasswordforserver
```

2. Generate the certificates again if you have a failure. Delete the files in the output directory before you rerun the `createSRACerts` script.
3. Stop all Storage Resource agents and the Data server.
For more information about starting or stopping IBM Spectrum Control services, see [Starting and stopping the IBM Spectrum Control servers](#).
4. Replace the certificate files:
 - Replace the certificate files for the Data server.
 - Replace the certificate files for the local Storage Resource agent that runs on the IBM Spectrum Control server.
 - Replace the certificate files for the remote Storage Resource agents that run on computers other than the IBM Spectrum Control server.
 - Replace the certificate files in the locations used for future installations of the remote Storage Resource agents.

Replace the certificate files for the Data server.

The new Data server certificate files are created in the following directory:

```
output_directory/sra_certs_out/server
```

By default, the `output_directory` is the directory where the `createSRACerts` script is run:

```
installation_dir/data/sra/tools/certs
```

These files are the Data server certificate files:

```
TPCDataServer.jks  
server.pwd
```

Copy the Data server certificate files to the following directory:

```
installation_dir/data/sra/certs
```

Replace the certificate files for the local Storage Resource agent that runs on the IBM Spectrum Control server.

The new Storage Resource agent certificates are created on the IBM Spectrum Control server in the following directory:

```
output_directory/sra_certs_out/agent
```

By default, the `output_directory` is the directory where the `createSRACerts` script is run:

```
installation_dir/data/sra/tools/certs
```

The Storage Resource agent certificate file is:

```
certs.zip
```

Copy the Storage Resource agent certificate file to the following directory on the IBM Spectrum Control server:

```
installation_dir/data/sra/server_operating_system
```

Where `server_operating_system` is the operating system on which the IBM Spectrum Control Data server is installed.

Extract the Storage Resource agent certificate file in the following directory on the IBM Spectrum Control server:

```
installation_dir/agent
```

Replace the certificate files for the remote Storage Resource agents that run on computers other than the IBM Spectrum Control server

The new Storage Resource agent certificates are created on the IBM Spectrum Control server in the following directory:

```
output_directory/sra_certs_out/agent
```

By default, the `output_directory` is the directory where the `createSRACerts` script is run:

```
installation_dir/data/sra/tools/certs
```

The Storage Resource agent certificate file is:

certs.zip

Copy the Storage Resource agent certificate file to the following directories on the IBM Spectrum Control server:

installation_dir/data/sra/remote_agent_operating_system

Where *remote_agent_operating_system* is an operating system on which a remote Storage Resource agent is installed.

Extract the Storage Resource agent certificate file in the following directory on the computer where the remote Storage Resource agent is installed:

installation_dir/agent

Replace the certificate files in the locations used for future installations of remote Storage Resource agents.

The new Storage Resource agent certificates are created on the IBM Spectrum Control server in the following directory:

output_directory/sra_certs_out/agent

By default, the **output_directory** is the directory where the **createSRACerts** script is run::

installation_dir/data/sra/tools/certs

The Storage Resource agent certificate file is:

certs.zip

Copy the Storage Resource agent certificate file to the following directories on the IBM Spectrum Control server:

installation_dir/data/sra/future_remote_agent_operating_system

Where *future_remote_agent_operating_system* is an operating system on which you install a Storage Resource agent some time in the future.

Restriction: This process assumes that the Storage Resource agent disk image can be modified. You must copy the installation files to a writable location before proceeding.

Before the Storage Resource agent can be installed locally, the new certificate must be copied to the agent system. Copy the new **certs.zip** Storage Resource agent certificate file from the **output_directory/sra_certs_out/agent** directory on the IBM Spectrum Control server to the agent system.

- a. On the agent system, extract the Storage Resource agent installation image in the **SRA_image_install_directory**.
- b. Copy the new **certs.zip** file into the following directory:

SRA_image_install_directory/sra/agent_operating_system

- c. Extract the new **certs.zip** file in the following directory:

SRA_image_install_directory/sra/agent_operating_system

Note: The **SRA_image_install_directory** value is the directory where the Storage Resource agent image was extracted and *agent_operating_system* is the directory that is named for the operating system that is running on the computer where you intend to install the Storage Resource agent.

- d. Install the Storage Resource agent with the wanted options.

5. Start the Data server and Storage Resource agents.

For more information about starting or stopping IBM Spectrum Control services, see [Starting and stopping the IBM Spectrum Control servers](#).

Related information

- [Installing Storage Resource agents](#)
- [Installing Storage Resource agents by using a command](#)

Configuration guidelines for 500 or more agents

You can use this information to help you manage 500 or more Storage Resource agents in IBM Spectrum® Control.

About this task

If you have 500 or more Storage Resource agents communicating with IBM Spectrum Control, complete the following steps:

Procedure

1. Probe the servers at least once a day or more, depending on when you want to test for alert conditions.
2. Set the following parameters in the **server.config** file:

MaxConnections=1200

The default is 500. Agents can have multiple connections to the server.

routerThreads=3 (max)

Incoming connections need to be routed to the correct Data Manager "service" queue and can stack up behind this thread. The server service runs the router and the agent service is where the connections are queued once routed and saved by any of three threads here to the repository.

3. Set the following parameter in the **Scheduler.config** file:

MaxSubmitthreads=8

Tells how many threads are used to tell the agents to start a job. Agent connections can queue up the scheduler service. After a job is run, the agent makes a connection to communicate with this thread to give it the job status.

Including a Storage Resource agent with a server golden image

If you use a golden operating system image to deploy new servers in your environment, you can include the Storage Resource agent on that golden image. The golden image enables the agents to start and register with the IBM Spectrum® Control server automatically upon deployment. This support applies only to Storage Resource agents running in daemon mode.

About this task

The default agent directory is:

- For Windows: `C:\Program Files\IBM\TPC\agent`
- For UNIX: `/opt/IBM/TPC/agent`

Follow these instructions to include the IBM Spectrum Control agent on a golden image.

Procedure

1. Install the Storage Resource agent in daemon mode on the golden image system.
2. Stop the Storage Resource agent on the golden image system.
For the Windows system: Click `Start > Settings > Control Panel > Administrative Tools > Services`. Stop the following service: IBM Spectrum Control Storage Resource Agent - *directory*. *directory* is where the Storage Resource agent is installed. The default directory is `installation_dir\agent`.

For the UNIX or Linux® system, run the following commands:

```
cd /opt/IBM/TPC/agent/bin/  
./agent.sh stop
```

3. Create one of the following files in the root directory for the agent. These files can be empty. Any content in these files is ignored.

REGISTERSRA

The file name must be uppercase with no file extension. This file causes the agent to run a probe and then register with the server. This file will use the existing Globally Unique Identifier (GUID).

REGISTERSRA_REGENGUID

The file name must be uppercase with no file extension. This file causes the agent to regenerate a new Globally Unique Identifier (GUID), run a probe, and then register with the server.

4. Delete the contents of the `agent_installation_directory/logs` directory. This clears any existing log messages so that you can view new messages that are logged.
5. Create the golden image copies of this system.
6. When a new system is preinstalled from this image and then started, the REGISTERSRA or REGISTERSRA_REGENGUID file is run. The Storage Resource agent automatically registers with the new IBM Spectrum Control server. You can then use the GUI to manage the Storage Resource agent deployment. For example, to confirm that the Storage Resource agent was deployed successfully, go to the Servers page and refresh the list.

Checking for a fully qualified host name

IBM Spectrum® Control requires fully qualified host names. Some machines might be configured to return a short host name, such as `system1` instead of a fully qualified host name, such as `system1.tpc.example.com`. This topic provides information on how to check for a fully qualified host name.

- [Checking for a fully qualified host name for AIX systems](#)
This topic provides information on how to verify a fully qualified host name for AIX®.
- [Checking for a fully qualified host name for Linux systems](#)
This topic provides information on how to verify a fully qualified host name for Linux®.
- [Checking for a fully qualified host name for Oracle Solaris](#)
This topic provides information about how to verify a fully qualified host name for Oracle Solaris systems.
- [Checking for a fully qualified host name for Windows systems](#)
Verify the fully qualified host name on Windows operating systems.

Checking for a fully qualified host name for AIX systems

This topic provides information on how to verify a fully qualified host name for AIX®.

About this task

The default domain name search order is as follows:

1. Domain Name System (DNS) server
2. Network Information Service (NIS)
3. Local `/etc/hosts` file.

If the `/etc/resolv.conf` file does not exist, the `/etc/hosts` file is used. If only the `/etc/hosts` file is used, the fully qualified computer name must be the first one that is listed after the IP address.

Verify that the `/etc/resolv.conf` file exists and contains the appropriate information, such as:

```
domain mydivision.mycompany.com
nameserver 123.123.123.123
```

If NIS is installed, the `/etc/irs.conf` file overrides the system default. It contains the following information:

```
hosts = bind,local
```

The `/etc/netsvc.conf` file, if it exists, overrides the `/etc/irs.conf` file and the system default. It contains the following information:

```
hosts = bind,local
```

If the NSORDER environment variable is set, it overrides all of the preceding files. It contains the following information:

```
export NSORDER=bind,local
```

Checking for a fully qualified host name for Linux systems

This topic provides information on how to verify a fully qualified host name for Linux®.

About this task

Linux uses a resolver library to obtain the IP address corresponding to a host name. The `/etc/host.conf` file specifies how names are resolved. The entries in the `/etc/host.conf` file tell the resolver library what services to use, and in what order, to resolve names. Edit the `host.conf` file using the `vi` editor to add the following lines:

```
# Lookup names through DNS first then fall back to /etc/hosts.
order bind,hosts
# Machines with multiple IP addresses.
multi on
# Check for IP address spoofing.
nospoof on
```

The **order** option indicates the order of services. The sample entry specifies that the resolver library should first consult the name server to resolve a name and then check the `/etc/hosts` file. It is recommended to set the resolver library to first check the name server, **bind** file, and then the **hosts** file (hosts) for better performance and security on all your servers. You must have the DNS and BIND software installed for this configuration to work.

The **multi** option determines whether a host in the `/etc/hosts` file can have multiple IP addresses. Hosts that have more than one IP address are said to be multihomed, because the presence of multiple IP addresses implies that the host has several network interfaces.

The **nospoof** option takes care of not permitting spoofing on this machine. IP-Spoofing is a security exploit that works by tricking computers into a trust relationship that you are someone that you really are not. In this type of attack, a machine is set up to look like a legitimate server and then issue connections and other types of network activities to legitimize end systems, other servers, or large data repository systems. This option must be set ON for all types of servers.

Checking for a fully qualified host name for Oracle Solaris

This topic provides information about how to verify a fully qualified host name for Oracle Solaris systems.

About this task

Verify that the `/etc/resolv.conf` file exists and contains the appropriate information, such as:

```
domain mydivision.mycompany.com
nameserver 123.123.123.123
```

A short name is used if the `/etc/nsswitch.conf` file contains a line that begins as follows and if the `/etc/hosts` file contains the short name for the computer:

```
hosts: files
```

To correct this problem, follow these steps:

Procedure

1. Change the line in the `/etc/nsswitch.conf` file to the following:

```
hosts: dns nis files
```

2. Enter the following command to stop the **inet** service:

```
/etc/init.d/inetd stop
```

3. Enter the following command to restart the **inet** service:

```
/etc/init.d/inetd start
```

Checking for a fully qualified host name for Windows systems

Verify the fully qualified host name on Windows operating systems.

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	a. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start. b. Click Control Panel, and then click System. c. Click Change Settings, click Change, and then click Change again.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	a. Click Start, Control Panel, System and Security. b. Click System, and then click Change Settings. c. On the Computer Name tab, click Change.

2. In the Computer name field, enter the fully qualified host name, and then click More.
3. Verify that the Primary DNS suffix field contains a domain name, and then click OK.

Granting local administrative privileges to a domain account

Automatically grant administrative privileges to Windows domain accounts. The user account for the Storage Resource agent requires local administrative rights. Because these rights are not necessarily guaranteed for domain users in a Windows domain environment, you are shown how to grant local administrative rights to domain users. Using this procedure, you do not have to manually process each machine in the domain.

About this task

Note: These steps are for a Windows system that is a member of a Windows domain and not for the Windows Domain Primary Domain Controller. To use Group Policy to grant local administrative privileges to a domain account, complete the following steps:

Procedure

1. On the domain controller, go to Administrative Tools, Active Directory Users and Computers (you must be running with Domain Administrator privileges).
2. Right-click on the Organizational Unit (OU) upon which you want to apply the Group Policy. Click Properties.
3. The Group Policy Properties panel is displayed. Select the Group Policy tab and click New to create a Group Policy.
4. Designate a name for the new Group Policy. Select the new Group Policy and click Edit.
5. The Group Policy Object Editor panel is displayed. Go to New Group Policy Object *your_policy*, Computer Configuration, Windows Settings, Security Settings, Restricted Groups. Right-click Restricted Groups, and then click Add Group.
6. For example, name the new group Administrators. Under Properties, add the user Administrator, and the domain accounts or groups upon which you want the Group Policy in effect for. For example, you can add TPC\storageadmin, TPC\storagegroup, and TPC\TestGroup. Click OK.
7. Add these user rights to the domain account:
 - Act as part of the operating system
 - Log on as a serviceIn the Group Policy Object Editor, go to New Group Policy Object *your_policy*, Computer Configuration, Windows Settings, Security Settings, Local Policies, User Rights Assignments. In the content pane, select "Log on as a service" and double-click. Add the domain user for whom you are granting user rights and click OK. Repeat this step for "Act as part of the operating system."
8. The group policy is now enforced for the Organizational Unit to include the domain accounts and groups specified under the local Administrators group on each computer in the Organizational Unit. In addition, the domain user has been granted the necessary rights. To verify the user rights, log in to a domain computer and open the Computer Management console. Select Groups, double-click the Administrators group, and verify the membership of the domain users.

Importing authentication information for a Storage Resource agent

The Storage Resource agent is installed as a non-daemon or daemon process. IBM Spectrum® Control stores the authentication information to connect to the host on which the Storage Resource agent has installed for the non-daemon agent. This authentication information can be changed depending on the environment.

To change the authentication information for a Storage Resource agent for non-daemon service, follow these steps:

1. Export the authentication information for a Storage Resource agent.
2. The data file exported contains information such as the host name, user ID, password, certificate location, and passphrase for every agent selected. The information is separated by the pipe character (|). For example,

```
agent_host|user|password|certificate|passphrase
```

You can update the password or passphrase in encrypted format or plain text format. If you want to update the password or passphrase in encrypted format, then you can use the **tpctool**. For example, go to this directory and run the **tpctool**:

```
cd installation_dir/cli
tpctool encrypt string_to_be_encrypted
```

This generates an encrypted string. Place this string in the data file to be imported and add @ENC@ to the end of the encrypted string. For example,

```
agent_host|usera|encrypted_password@ENC@|certificate|
encrypted_passphrase@ENC@
```

encrypted_password is the encrypted string for the password and *encrypted_passphrase* is the encrypted string for the passphrase.

3. Import the data file.

Installing and configuring the IBM Spectrum Control server with multiple NIC cards

If your IBM Spectrum® Control server has multiple network interface cards (NIC), install the IBM Spectrum Control server using a fully qualified hostname that resolves to the IP address of NIC card you want to use. After you install the server, all incoming and outgoing communication are successfully handled.

Installing IBM Spectrum Control for a multiple network configuration

If the IBM Spectrum Control server you are installing has multiple NIC, and is configured to use multiple network addresses, ensure that you use the fully qualified hostname that resolves to the appropriate IP address during installation. You can either setup the HOSTS file or the DNS to resolve the fully qualified host names to appropriate IP addresses.

Outgoing communication initiated by the IBM Spectrum Control server

All the outgoing communication that is initiated by the IBM Spectrum Control server is not affected if the server is configured for a multiple network environment.

For example, if you have a IBM Spectrum Control server with two IP addresses: 10.10.10.11 and 9.9.9.10, and 10.10.10.11 is used during installation, all outgoing transmissions can be sent to the devices and agents in both networks.

The following list includes examples of outgoing communication that is initiated by the IBM Spectrum Control server:

- Storage systems using native interfaces
 - Run probe, performance monitor, and provisioning. Examples of storage systems that use the native interface include SAN Volume Controller, Storwize® V7000 Unified, Storwize V7000, and XIV® system.
- Switches (SNMP and SMIS-S providers)
 - Run SNMP and SMI-S provider probes.
- SMI-S providers
 - IBM Spectrum Control uses SMI-S providers (CIM agents) for the managed objects to gather information about the resource.
- VMware vCenter
 - Run probes.
- Agents (Storage Resource agents)
 - Deploy agents, run data collection, and run scripts.
- IBM Spectrum Control servers
 - Run data collection.

Incoming communication that is initiated by the resources, agents, and GUI

Incoming communication that is initiated by the resources or agents can work with only the IP address that is specified during the installation with the exception of DS8000® events.

For DS8000 events, the IBM Spectrum Control server must initiate and establish a socket connection directly with the Hardware Management Console (HMC) to receive events. The DS8000 HMC uses that socket connection to send events. As long as the IBM Spectrum Control server can initiate the communication to the HMC, DS8000 events can be received.

IBM Spectrum Control informs resources and agents to initiate communication to the IP address provided during the installation. This example uses the IP address 10.10.10.11. However, depending on the communication, you might be able to change the IP address. For example, IBM Spectrum Control does not configure SAN switches to send SNMP traps to IBM Spectrum Control, so you can use either 9.9.9.10 or 10.10.10.11.

The following list includes examples of incoming communication that are initiated by the resources, agents, and the GUI:

- DS8000 events
 - Events sent by the HMC to the IBM Spectrum Control server
- SNMP trap notifications
 - SNMP traps sent from the switches and other resources
- CIM indications
 - Indications sent by the SMI-S providers (CIM agents).
- Servers (agents)
 - Job results and registration
- IBM Spectrum Control GUI
 - Any request.

CIM indications

A CIM indication is an event that occurs on a managed object, for example, the completion or failure of an operation. The CIM indications are managed by the CIM object manager. IBM Spectrum Control uses the SMI-S providers for the managed objects to gather information about the resource.

Manually customize CIM indications on a IBM Spectrum Control system that has multiple IP addresses. To configure IBM Spectrum Control to receive CIM indications in an IPv4, IPv6, and dual stack (IPv4 and IPv6) environment, see [Configuring IBM Spectrum Control with multiple IP addresses](#).

The manual customization task does not apply to storage devices that use the native interfaces.

Replacing the default SSL certificate for the Device, Alert, or Web server with a self-signed certificate

To replace the default SSL certificate for the Device, Alert, or Web server, with a self signed certificate, use the IBM® Key Management (iKeyman) utility.

About this task

If you have strong security requirements, you might want to replace the default certificate for the Web server, with a self-signed certificate so that you can securely connect to the Web server while you use the `https` protocol.

Procedure

1. Log on to the server where IBM Spectrum® Control is installed. Ensure that you log on with the appropriate user privileges.
2. Open the `/jre/bin` directory where IBM Spectrum Control is installed.
3. Enter the iKeyman utility command.
For Windows operating system, enter the following command:

```
ikeyman.exe
```


For AIX® or Linux® operating system, enter the following command:

```
./ikeyman
```
4. Click Key Database File > Open.
5. Complete the following tasks:
 - a. Set the Key database type to PKCS12.
 - b. In the **File Name** field, click Browse. Note: You can replace only one certificate at a time.
To replace the default SSL certificate for the Device server, go to the `installation_dir/wlp/usr/servers/deviceServer/resources/security/` directory, select the `key.p12` file, and click Open.

To replace the default SSL certificate for the Alert server, go to the `installation_dir/wlp/usr/servers/alertServer/resources/security/` directory, select the `key.p12` file, and click Open.

To replace the default SSL certificate for the Web server, go to the `installation_dir/wlp/usr/servers/webServer/resources/security/` directory, select the `key.p12` file, and click Open.
 - c. Click OK.
6. On the Password Prompt page, type `default`, and click OK.
The Personal Certificates list contains only the certificate with the `default` label.
To replace the default certificate with a new self-signed certificate, complete the following tasks:
 - a. Click New Self-Signed.
 - b. On the Create New Self-Signed Certificate page, enter a unique value in the Key Label field.
 - c. Provide values for the other fields, and click OK.
The list of Personal Certificates contains your new self-signed certificate with the label that you provided and the old self signed certificate with the `default` label.
 - d. Select the old self signed certificate with the `default` label and click Rename.
 - e. Enter a new label for the old self signed certificate, and click OK.
 - f. Select the new self signed certificate and click Rename.
 - g. Enter `default` as the new label, for the new self-signed certificate, and click OK.
7. In the iKeyman utility, click Key Database File > Exit.
8. Stop and start the Device, Alert, or Web server.

Related tasks

- [Updating IBM Spectrum Control data collector trusted certificates after replacing default SSL certificate for the Device server with a self-signed certificate OR an external certificate](#)
- [Starting and stopping the IBM Spectrum Control servers](#)

Replacing the default SSL certificate for the Device, Alert, or Web server with a certificate from an external certificate authority

To replace the default SSL certificate for the Device, Alert, or Web server, with a certificate from an external certificate authority, use the IBM® Key Management (iKeyman) utility.

About this task

If you have strong security requirements, you might want to replace the default certificate for the Web server, with a certificate from an external certificate authority so that you can securely connect to the Web server while you use the `https` protocol. When you replace the existing certificate, it can remove web browser certificate error warnings.

Procedure

1. Log on to the server where IBM Spectrum® Control is installed. Ensure that you log on with the appropriate user privileges.
2. Open the `/jre/bin` directory where IBM Spectrum Control is installed.
3. Enter the iKeyman utility command.
For Windows operating system, enter the following command:

```
ikeyman.exe
```


For AIX® or Linux® operating system, enter the following command:

```
./ikeyman
```

4. Click Key Database File > Open.
 5. Complete the following tasks:
 - a. Set the Key database type to PKCS12.
 - b. In the **File Name** field, click Browse. Note: You can replace only one certificate at a time.
 To replace the default SSL certificate for the Device server, go to the `installation_dir/wlp/usr/servers/deviceServer/resources/security/` directory, select the key.p12 file, and click Open.

 To replace the default SSL certificate for the Alert server, go to the `installation_dir/wlp/usr/servers/alertServer/resources/security/` directory, select the key.p12 file, and click Open.

 To replace the default SSL certificate for the Web server, go to the `installation_dir/wlp/usr/servers/webServer/resources/security/` directory, select the key.p12 file, and click Open.
 - c. Click OK.
 6. In the iKeyman utility, select Create > New Certificate Request.
 7. Enter a unique value in the Key Label field and provide values for the other fields.
 8. Pay special attention to the value you provide in the Enter the name of a file in which to store the certificate request field and click OK.
 A message is displayed that informs you the location of the file that contains your new certificate request. You need to send the new certificate request file to your external certificate authority.
 9. On the Message page, click OK.
 The external certificate authority signs your new certificate request and sends back your new certificate. The external certificate authority might send their signer certificate or the external certificate authority might assume that you already have their signer certificate in the key database file.
 If the external certificate authority sends their signer certificate, complete the following tasks:
 - a. Select Signer Certificates and click Add.
 - b. Provide the File Name and Location values of the file that contains the Signer Certificate and click OK.
 - c. Enter a label for the signer certificate, and click OK.
 If the external certificate authority assumes that you already have their signer certificate in the key database file, complete the following tasks:
 - d. Select Signer Certificates and click Populate.
 - e. Search the lists of CA Certificates, select the one or the ones for the external certificate authority that signed your new certificate request, and click OK.
 If the lists of CA Certificates do not contain the one(s) for the external certificate authority that signed your new certificate request, ask your external certificate authority to send their signer certificate.
 After you have the signer certificate for the external certificate authority in the keystore, complete the following tasks to receive the new certificate signed by the external certificate authority:
 - f. Select Personal Certificates and click Receive.
 - g. Provide the File Name and Location values of the file that contains your new certificate from the external certificate authority and click OK.
 - h. Select the old self-signed certificate with the `default` label and click Rename.
 - i. Enter a new label for the old self-signed certificate and click OK.
 - j. Select your new certificate from the external certificate authority and click Rename.
 - k. Enter `default` as the new label for the new certificate from the external certificate authority and click OK.
 10. In the iKeyman utility, click Key Database File > Exit.
 11. Stop and start the Device, Alert, or Web server.
- [Updating IBM Spectrum Control data collector trusted certificates after replacing default SSL certificate for the Device server with a self-signed certificate OR an external certificate](#)
 Use the **keytool** command to update the IBM Spectrum Control data collector trusted certificates after you replace the default SSL certificate for the IBM Spectrum Control Device server.

Related tasks

- [Updating IBM Spectrum Control data collector trusted certificates after replacing default SSL certificate for the Device server with a self-signed certificate OR an external certificate](#)
- [Starting and stopping the IBM Spectrum Control servers](#)

Updating IBM Spectrum Control data collector trusted certificates after replacing default SSL certificate for the Device server with a self-signed certificate OR an external certificate

Use the **keytool** command to update the IBM Spectrum® Control data collector trusted certificates after you replace the default SSL certificate for the IBM Spectrum Control Device server.

About this task

If you replace the default SSL certificate for the IBM Spectrum Control Device server, you must update the IBM Spectrum Control data collector trusted certificates or else the data collector does not communicate properly with the Device server.

Procedure

1. Log on to the server where IBM Spectrum Control is installed. Ensure that you log on with the appropriate user privileges.
2. Stop the Device server.
3. Open the `/jre/bin` directory where IBM Spectrum Control is installed.
4. Enter the following command to export the default SSL certificate from the Device server keystore.
 For Windows operating system, enter the following command:

```
keytool.exe -exportcert -alias default  
-keystore "installation_dir\wlp\usr\servers\deviceServer\resources\security\key.p12" -storetype pkcs12  
-storepass device_server_keystore_password -file deviceServer.cert
```

Where *device_server_keystore_password* is the Device server keystore password and the default value for this password is *default*.

For AIX® or Linux® operating system, enter the following command:

```
./keytool -exportcert -alias default  
-keystore installation_dir/wlp/usr/servers/deviceServer/resources/security/key.p12 -storetype pkcs12  
-storepass device_server_keystore_password -file deviceServer.cert
```

Where *device_server_keystore_password* is the Device server keystore password and the default value for this password is *default*.

5. Enter the following command to delete the previous IBM Spectrum Control Device server SSL certificate from the IBM Spectrum Control data collector trusted certificates.

For Windows operating system, enter the following command:

```
keytool.exe -delete -alias deviceServer -keystore "installation_dir\jre\lib\security\cacerts" -storepass  
data_collector_keystore_password
```

Where *data_collector_keystore_password* is the IBM Spectrum Control data collector keystore password and the default value for this password is *changeit*.

For AIX or Linux operating system, enter the following command:

```
./keytool -delete -alias deviceServer -keystore installation_dir/jre/lib/security/cacerts  
-storepass data_collector_keystore_password
```

Where *data_collector_keystore_password* is the IBM Spectrum Control data collector keystore password and the default value for this password is *changeit*.

6. Enter the following command to add the default SSL certificate from the IBM Spectrum Control Device server to the IBM Spectrum Control data collector trusted certificates.

For Windows operating system, enter the following command:

```
keytool.exe -importcert -noprompt -trustcacerts -alias deviceServer -file deviceServer.cert  
-keystore "installation_dir\jre\lib\security\cacerts" -storepass data_collector_keystore_password
```

Where *data_collector_keystore_password* is the IBM Spectrum Control data collector keystore password and the default value for this password is *changeit*.

For AIX or Linux operating system, enter the following command:

```
./keytool -importcert -noprompt -trustcacerts -alias deviceServer -file ./deviceServer.cert  
-keystore installation_dir/jre/lib/security/cacerts -storepass data_collector_keystore_password
```

Where *data_collector_keystore_password* is the IBM Spectrum Control data collector keystore password and the default value for this password is *changeit*.

7. Start the Device server.

Related tasks

- [Starting and stopping the IBM Spectrum Control servers](#)

Replacing the default SSL certificate for the Export server

You can replace the default SSL certificate for the Export server by adding the certificate and private key file to the appropriate directory and restarting the Export server.

About this task

The Export server is used by the Web server to generate certain types of reports. By default, the certificate for the Export server is self-signed and the corresponding private key does not have a password. However, if you have strong security requirements, you might want to replace the default certificate and the corresponding private key used by the Export server.

Procedure

1. Stop the Export server.
2. If your signed certificate has an accompanying chain certificate, append the contents of the chain certificate to the bottom of your signed certificate file.
3. Place the signed certificate in the *installation_dir/export/conf/* directory.
4. Place the matching private key file in the *installation_dir/export/conf/* directory.
5. If your private key file requires a password, create and place the password in the *installation_dir/export/conf/* directory.
6. Start the Export server again.

Related tasks

- [Starting and stopping the IBM Spectrum Control servers](#)

Generating a new default self-signed SSL certificate for the Export server

You can generate a new, default self-signed SSL certificate for the Export server by using the **openssl** command.

Before you begin

You must have the extracted IBM Spectrum® Control installation image files present on the IBM Spectrum Control server where you are going to generate a new, default self-signed SSL certificate for the Export server.

Procedure

1. Log on to the server where IBM Spectrum Control is installed.
Ensure that you log on with the appropriate user privileges.
2. Stop the Export server.
3. Enter the **openssl** command to generate a new, default self-signed SSL certificate for the Export server.
For the Windows operating system, enter the following command:

```
installation_dir\data\sra\tools\openssl\openssl req  
-config extracted_image_dir\scripts\export\openssl.cfg  
-new -newkey rsa:2048 -x509 -nodes -keyout installation_dir\export\conf\export.key  
-out installation_dir\export\conf\export.cert -days 3650  
-subj /C=us/O=ibm/OU=exportServer/CN=machine_FQDN
```

Where *installation_dir* is the location where IBM Spectrum Control is installed, *extracted_image_dir* is the location where the IBM Spectrum Control installation image is extracted, and *machine_FQDN* is the fully qualified domain name of the machine where you installed IBM Spectrum Control. For example, *myserver.mycompany.com*.

If your *installation_dir* or *extracted_image_dir* location contains spaces, use double quotes around those paths.

For example:

```
"C:\Program Files\IBM\TPC\data\sra\tools\openssl\openssl" req  
-config C:\Downloads\SC-Image\SC\scripts\export\openssl.cfg  
-new -newkey rsa:2048 -x509 -nodes -keyout "C:\Program Files\IBM\TPC\export\conf\export.key"  
-out "C:\Program Files\IBM\TPC\export\conf\export.cert" -days 3650  
-subj /C=us/O=ibm/OU=exportServer/CN=myserver.mycompany.com
```

For the AIX® or Linux® operating systems, enter the following command:

```
installation_dir/data/sra/tools/openssl/openssl req  
-config extracted_image_dir/scripts/export/openssl.cfg  
-new -newkey rsa:2048 -x509 -nodes -keyout installation_dir/export/conf/export.key  
-out installation_dir/export/conf/export.cert -days 3650  
-subj /C=us/O=ibm/OU=exportServer/CN=machine_FQDN
```

Where *installation_dir* is the location where IBM Spectrum Control is installed, *extracted_image_dir* is the location where the IBM Spectrum Control installation image is extracted, and *machine_FQDN* is the fully qualified domain name of the machine where you installed IBM Spectrum Control. For example, *myserver.mycompany.com*.

4. Start the Export server.

Related tasks

- [Starting and stopping the IBM Spectrum Control servers](#)

Enabling TLS 1.0 and 1.1 for ports

IBM Spectrum® Control uses Transport Layer Security (TLS) to secure communications between IBM Spectrum Control components.

Restriction: IBM Spectrum Control uses TLS 1.2 protocol for communicating on ports. It does not use TLS 1.3, and TLS 1.1 and 1.0 are disabled by default for increased security.

- [Enabling TLS 1.1 and 1.0 for IBM Spectrum Control ports](#)
To enable TLS 1.1 and 1.0 for IBM Spectrum Control ports, update the `java.security` file (Alert, Data, Device, and Web server) and the `server.config` file (Export server).

Enabling TLS 1.1 and 1.0 for IBM Spectrum Control ports

To enable TLS 1.1 and 1.0 for IBM Spectrum Control ports, update the `java.security` file (Alert, Data, Device, and Web server) and the `server.config` file (Export server).

Before you begin

IBM Spectrum Control ports have TLS 1.1 and 1.0 disabled by default for increased security. Therefore, IBM Spectrum Control will not be able to communicate with resources that do not support TLS 1.2. If you want to upgrade your resources to a version that supports TLS 1.2, contact your vendor. You can also re-enable TLS 1.1 and 1.0 for IBM Spectrum Control ports.

Procedure

1. Stop all IBM Spectrum Control servers.
2. Open the `installation_dir/jre/lib/security/java.security` file.
3. To enable TLS 1.1 and 1.2 in the Alert, Data, Device, and Web server, remove the "`TLSv1.1` ," text from the `jdk.tls.disabledAlgorithms` line.
BEFORE:

```
jdk.tls.disabledAlgorithms=MD5withRSA, DH keySize < 1024, TLSv1, TLSv1.1 , EC keySize < 224, anon, NULL
```

AFTER:

```
jdk.tls.disabledAlgorithms=MD5withRSA, DH keySize < 1024, TLSv1, EC keySize < 224, anon, NULL
```

To enable TLS 1.0, 1.1, and 1.2 in the Alert, Data, Device, and Web server, remove the "TLSv1, TLSv1.1 , " text from the `jdk.tls.disabledAlgorithms` line.

BEFORE:

```
jdk.tls.disabledAlgorithms=MD5withRSA, DH keySize < 1024, TLSv1, TLSv1.1 , EC keySize < 224, anon, NULL
```

AFTER:

```
jdk.tls.disabledAlgorithms=MD5withRSA, DH keySize < 1024, EC keySize < 224, anon, NULL
```

4. Open the `installation_dir/export/conf/server.config` file.

5. To enable only TLS 1.1 in the Export server, change the "secureProtocol" value from "TLSv1_2_method" to "TLSv1_1_method".

For example:

```
"secureProtocol": "TLSv1_1_method"
```

To enable only TLS 1.0 in the Export server, change the "secureProtocol" value from "TLSv1_2_method" to "TLSv1_method".

For example:

```
"secureProtocol": "TLSv1_method"
```

To enable TLS 1.0, 1.1, and 1.2 in the Export server, change the "secureProtocol" value from "TLSv1_2_method" to "".

For example:

```
"secureProtocol": ""
```

You cannot configure the Export server such that only TLS 1.1 and 1.2 are enabled.

6. Restart the IBM Spectrum Control servers.

Configuring Db2, AIX, and Linux for IPv6-only environment

Use this information to configure Db2®, AIX®, and Linux® for an IPv6-only environment.

Configuring the AIX system for IPv6 only

For IPv6 support, the AIX operating system must have level TL 5300–06 installed.

To configure the AIX operating system for IPv6, complete the following steps:

1. Obtain the most recent versions of **openssh** and **openssl** packages for AIX and install them. Some older version of **openssh** does not work in an IPv6-only environment.
2. Change **sshd** (Secure Shell Daemon) on AIX system to accept IPv6 connections.
 - a. In the `/etc/ssh/sshd_config` file, uncomment the line "ListenAddress:".
 - b. Restart **sshd** with the following commands:

```
stopsrc -g ssh
startsrc -g ssh
```
 - c. From another IPv6 system, verify that you contact AIX over IPv6 (by using ssh).
3. In SMIT, set the IPv4 address to 0.0.0.0 for all interfaces. Save the file.
4. Edit the `/etc/resolv.conf` file to use IPv6 DNS server or servers.

Configuring Db2 on AIX for IPv6 systems

To get Db2 on AIX operating systems to work on IPv6 systems, complete the following steps:

1. Identify the host name that is used by Db2 in the `db2nodes.cfg` file:

```
# cat ~db2inst1/sqllib/db2nodes.cfg
0 myhost 0
#
```

2. Edit the `/etc/hosts` file and make sure that the host name found in the `db2nodes.cfg` file resolves to an IPv6 address. Use the **vi** editor to verify that the host name is not on any line with an IPv4 address. In particular, ensure that the host name is not listed as an alias for the IPv4 loopback address 127.0.0.1.

```
# vi /etc/hosts
127.0.0.1 loopback localhost
::1 localhost
2001:db8:0:0:209:6bff:fe09:63fa myhost.mydomain myhost
```

3. Stop Db2 and set Db2 to use IPv6 addressing. Restart Db2.
 - a. Source the Db2 profile:

```
. ~db2inst1/sqllib/db2profile
```

- b. Stop Db2:

```
db2stop
```

c. Configure Db2 to use IPv6.

```
db2set
```

An example of the output is: **DB2FCMCOMM=TCPIP6**.
d. Start Db2.

```
db2start
```

In some installations, the AIX server does not have a graphical console that is attached to the server. In this situation, you can select another system with an X11 server to display the IBM Spectrum® Control installation and IBM Spectrum Control application. The X11 server must have IPv6 configured and an SSH client installed. Open an SSH connection from a shell on the X11 server desktop with the -X option to permit forwarding of X11 applications from the remote AIX server. Start the IBM Spectrum Control installation program or application from the SSH shell.

```
ssh -X my_IPv6_host  
/opt/IBM/TPC/gui/TPCD.sh
```

Configuring Db2 on Linux for IPv6-only systems

To get Db2 on Linux systems to work in an IPv6-only environment, follow these steps:

1. Install Db2 in dual-stack configuration.
2. Stop Db2 and set Db2 to use IPv6 addressing:
 - a. As the root user from the Linux command-line, run this command:

```
su - db2inst1
```

b. Stop Db2 by running this command:

```
db2stop
```

c. Configure Db2 to use IPv6 by running this command:

```
db2set
```

An example of the output is: **DB2FCMCOMM=TCPIP6**.

The host name in the **db2nodes.cfg** file resolves to an IPv6 address. This action can require you to change the domain or search directive in the **/etc/resolv.conf** file to specify a domain in which the host name can resolve to IPv6. You can also edit the **/etc/hosts** file so that the host name resolves to an IPv6 address.

d. Start Db2 by running this command:

```
db2start
```

Administering

Administer IBM Spectrum® Control and its components to ensure that your storage environment is being monitored as intended. Some administering tasks include stopping and starting product services, increasing memory allocation, monitoring the health of product components, and managing storage resources and data sources. You can use the Db2® command-line interface or IBM® Data Studio to administer Db2.

- [Administering resources and data sources](#)
Administer monitored resources and the data sources that are associated with those resources. Data sources can be agents that manage resources or VMware vCenter servers. An agent might be a CIM agent or a Storage Resource agent.
- [Starting and stopping the IBM Spectrum Control servers](#)
You can start and stop the IBM Spectrum Control servers in the GUI or by running scripts.
- [Checking the version and license of IBM Spectrum Control](#)
The version and license of IBM Spectrum Control that is installed on your system determine the IBM Spectrum Control functions that are available.
- [Checking IBM Spectrum Control status](#)
The System Management page shows a high-level summary of the status of the server or servers on which IBM Spectrum Control is installed. Use the System Management page to troubleshoot problems with the system, create trace logs, and get technical support.
- [Increasing the memory allocation for the Data server](#)
If the data memory that is allocated for your Data server is insufficient, you can increase the memory. The default maximum memory for the Data server is set to 1024 MB.
- [Changing passwords](#)
IBM Spectrum Control provides a GUI and non-GUI password tool; however, both tools achieve the same purpose.
- [Granting local administrative privileges to a domain account](#)
Automatically grant administrative privileges to Windows domain accounts. The user account for the Storage Resource agent requires local administrative rights. Because these rights are not necessarily guaranteed for domain users in a Windows domain environment, you are shown how to grant local administrative rights to domain users. Using this procedure, you do not have to manually process each machine in the domain.
- [Collecting diagnostic information about IBM Spectrum Control](#)
You can use the service tool to collect diagnostic information about IBM Spectrum Control. The tool detects the system configuration, collects the applicable information, and creates a compressed file that can be sent to IBM Software Support.
- [Administering the IBM Spectrum Control database](#)
The IBM Spectrum Control database is the repository for information that is collected about the monitored resources in your environment.
- [Administering Db2](#)
Administer IBM Db2 by backing up the IBM Spectrum Control database, starting the IBM Data Studio full client, and starting and stopping Db2.

Administering resources and data sources

Administer monitored resources and the data sources that are associated with those resources. Data sources can be agents that manage resources or VMware vCenter servers. An agent might be a CIM agent or a Storage Resource agent.

- **[Storage systems](#)**
Administer the storage systems that are monitored by IBM Spectrum Control. Administering actions include adding and removing storage systems, updating credentials, and testing connections.
- **[Hypervisors and VMware data sources](#)**
Administer the hypervisors, vCenter Server Appliance systems and vCenter Server systems that are monitored by IBM Spectrum Control. vCenter Server Appliance systems and vCenter Server systems are data sources that can monitor multiple hypervisors. A hypervisor can be an ESX or ESXi host. Each hypervisor can host multiple virtual machines.
- **[Switches and fabrics](#)**
Administer the switches and fabrics that are monitored by IBM Spectrum Control. Administering actions include adding and removing switches and fabrics, modifying connection information, and testing connections.
- **[Servers and Storage Resource agents](#)**
Administer servers and the Storage Resource agents that collect asset, status, and file system information about servers.
- **[SMI-S providers](#)**
Administer SMI-S providers that are associated with storage resources that are monitored by IBM Spectrum Control. SMI-S providers enable communication between IBM Spectrum Control and certain types of storage systems and switches.
- **[SNMP agents](#)**
SNMP agents are switches and directors that communicate with IBM Spectrum Control through SNMP. [»](#)IBM Spectrum Control supports SNMPv1, SNMPv2, and SNMPv3 for connecting to switches and sending alert notifications to SNMP trap destinations. [«](#)

Storage systems

Administer the storage systems that are monitored by IBM Spectrum® Control. Administering actions include adding and removing storage systems, updating credentials, and testing connections.

- **[Viewing information about storage systems](#)**
View detailed information about storage systems that are monitored by IBM Spectrum Control.
- **[Updating the credentials for storage systems](#)**
Change the credentials that IBM Spectrum Control uses to authenticate to a storage system or the CIM agent that manages a storage system. You can also change the host name or the IP address.
- **[Testing the connection to a storage system](#)**
Verify that IBM Spectrum Control can communicate with a monitored storage system. For storage systems that are managed by a CIM agent or Storage Resource agent, the connection to the agent is tested.
- **[Collecting CIM agent logs](#)**
You can collect logs for certain IBM® CIM agents using the command line interface.
- **[Removing storage systems](#)**
Remove storage systems that you no longer want to monitor with IBM Spectrum Control.

Viewing information about storage systems

View detailed information about storage systems that are monitored by IBM Spectrum® Control.

Procedure

To view information about storage systems, complete the following steps:

1. Depending on how the storage system is configured, go to **Storage > Block Storage Systems**, **Storage > File Storage Systems**, or **Storage > Object Storage Systems**. Storage systems can be configured in the following ways:
 - Block storage system
Storage systems that are configured for storing or retrieving data only in block format include System Storage® DS series, SAN Volume Controller, Storwize® V7000, and other SAN-based storage systems.
Storage systems that can be configured for both file and block data include Storwize V7000 Unified and NetApp Filers.
 - File storage system
Storage systems that can be configured for both file and block data include Storwize V7000 Unified and NetApp Filers.
 - Object storage system
The storage system that can be configured for both file and object data is IBM Spectrum Scale storage system.
The storage system that can be configured only for object data is IBM® Cloud Object Storage.
- Information about monitored storage systems is displayed.
2. Right-click a storage system and select **View Properties** to view the key properties for the system.

Updating the credentials for storage systems

Change the credentials that IBM Spectrum® Control uses to authenticate to a storage system or the CIM agent that manages a storage system. You can also change the host name or the IP address.

Before you begin

If the storage system is managed by multiple data sources, for example multiple CIM agents, the menu is displayed as **Connections > Modify Connection > data_sources**. Select the data source for which you want to update the credentials.

The type of storage system determines the credentials that you can update.

- [Updating the credentials for a System Storage DS8000 storage system](#)
Change the credentials that IBM Spectrum Control uses to authenticate to a System Storage® DS8000® storage system.
- [Updating the credentials for an XIV or IBM Spectrum Accelerate](#)
Change the credentials that IBM Spectrum Control uses to authenticate to an XIV® or IBM Spectrum Accelerate.
- [Updating the credentials for storage systems that run IBM Spectrum Virtualize](#)
Change the credentials that IBM Spectrum Control uses to authenticate to IBM Spectrum Virtualize storage systems.
- [Updating the credentials for a Storwize V7000 Unified storage system](#)
Change the credentials that IBM Spectrum Control uses to authenticate to a Storwize® V7000 Unified storage system.
- [Updating the credentials for a storage system that is managed by a CIM agent](#)
Change the credentials that IBM Spectrum Control uses to authenticate to a CIM agent.

Updating the credentials for a System Storage DS8000 storage system

Change the credentials that IBM Spectrum® Control uses to authenticate to a System Storage® DS8000® storage system.

About this task

To update the credentials for a System Storage DS8000 storage system, complete the following steps. You can update the host name or IP address for the secondary HMC that is used to manage the storage system, the user name, and the password.

Procedure

1. In the menu bar, go to **Storage > Block Storage Systems**.
Information about monitored storage systems is displayed.
2. Right-click a storage system and click **Connections > Modify Connection**.
3. Change the secondary HMC host name or IP address, the user name, or the password, and then click OK.

Updating the credentials for an XIV or IBM Spectrum Accelerate

Change the credentials that IBM Spectrum® Control uses to authenticate to an XIV® or IBM Spectrum Accelerate.

About this task

To update the credentials for an XIV or IBM Spectrum Accelerate, complete the following steps. You can update the IP address or host name, the user name, and the password.

Procedure

1. In the menu bar, go to **Storage > Block Storage Systems**.
2. Right-click a storage system and click **Connections > Modify Connection**.
3. Change the host name or IP address, user name, or password, then click OK.

Updating the credentials for storage systems that run IBM Spectrum Virtualize

Change the credentials that IBM Spectrum® Control uses to authenticate to IBM Spectrum Virtualize storage systems.

Before you begin

In this documentation, IBM Spectrum Virtualize is used to refer collectively to IBM® SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, and IBM Storwize® storage systems, and to IBM FlashSystem® devices that run IBM Spectrum Virtualize.

About this task

You can update the IP address or host name, the user name, and the password.

Procedure

1. In the menu bar, go to **Storage > Block Storage Systems**.
Information about monitored storage systems is displayed.
2. Right-click a storage system and click **Connections > Modify Connection**.
3. Update the following credentials as required and then click OK:

Authentication

You can use a user name and password or a private Secure Shell (SSH) key to log on to the storage system. The authentication method that you select determines the options that are displayed.

User name/Password

The user name and password for logging on to the storage system.

Secure Shell (SSH)

Use a private SSH key to log in to the storage system.

Authenticate with an SSH key

When you upload the SSH private key, you must add the SSH user who is associated with that SSH private key when it was created.

SSH user

The user that was associated with the SSH key when it was created. The user must have an Administrator role so that configuration, capacity, status, and performance metadata can be collected about the device.

SSH password

The password associated with the SSH user. If the password was not created, leave this field blank.

SSH key

Upload the new SSH private key that was generated to authenticate with the storage system. The valid file formats for SSH keys are PEM and PuTTY. If the key is in another format such as OpenSSH, it must be converted to PEM or PuTTY before it can be used.

The SSH key file is transferred from the computer where the web browser is located to the computer where the IBM Spectrum Control server is located.

The default location is \${device.conf}\tpc_svc.pem, which represents the IBM Spectrum Control default key file tpc_svc.pem. The tpc_svc.pem file is in the conf directory where the Device server is installed.

Restriction: PuTTY Private Key version 3 files are not supported. If you use PuTTYgen 0.75 or later to generate keys, ensure that you choose the option to generate version 2 keys. In PuTTYgen, click 2 for PPK file version on the Private Key File Parameters window before you generate the key.

For more information about PuTTYgen, see <https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html>.

Passphrase

The passphrase that is associated with the SSH key pair. If a passphrase was not created for the SSH key pair, leave the field blank. For SSH keys that use the PuTTY file format, you cannot use a passphrase to protect the private key.

Updating the credentials for a Storwize V7000 Unified storage system

Change the credentials that IBM Spectrum® Control uses to authenticate to a Storwize® V7000 Unified storage system.

About this task

To update the credentials for a Storwize V7000 Unified storage system, complete the following steps. You can update the IP address or host name, the user name, and the password.

Procedure

1. In the menu bar, go to Storage > Block Storage Systems.
Information about monitored storage systems is displayed.
2. Right-click a storage system and click Connections > Modify Connection.
3. Update the following credentials as required, and then click OK:

Authentication

You can use a user name and password or a private Secure Shell (SSH) key to log on to the storage system. The authentication method that you select determines the options that are displayed.

User name/Password

The user name and password for logging on to the storage system.

Secure Shell (SSH)

Use a private SSH key to log in to the storage system.

Authenticate with an SSH key

When you upload the SSH private key, you must add the SSH user who is associated with that SSH private key when it was created.

SSH user

The user that was associated with the SSH key when it was created. The user must have an Administrator role so that configuration, capacity, status, and performance metadata can be collected about the device.

SSH password

The password associated with the SSH user. If the password was not created, leave this field blank.

SSH key

Upload the new SSH private key that was generated to authenticate with the storage system. The valid file formats for SSH keys are PEM and PuTTY. If the key is in another format such as OpenSSH, it must be converted to PEM or PuTTY before it can be used.

The SSH key file is transferred from the computer where the web browser is located to the computer where the IBM Spectrum Control server is located.

The default location is \${device.conf}\tpc_svc.pem, which represents the IBM Spectrum Control default key file tpc_svc.pem. The tpc_svc.pem file is in the conf directory where the Device server is installed.

Restriction: PuTTY Private Key version 3 files are not supported. If you use PuTTYgen 0.75 or later to generate keys, ensure that you choose the option to generate version 2 keys. In PuTTYgen, click 2 for PPK file version on the Private Key File Parameters window before you generate the key.

For more information about PuTTYgen, see <https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html>.

Passphrase

The passphrase that is associated with the SSH key pair. If a passphrase was not created for the SSH key pair, leave the field blank. For SSH keys that use the PuTTY file format, you cannot use a passphrase to protect the private key.

Use different authentication credentials for file storage

Storwize V7000 Unified contains block-level and file-level data. If the credentials are different for block storage and file storage, select this check box to define the credentials for file storage. The options and fields that are displayed are described previously under Authentication.

Tip: If you use an SSH key to log on to the file module, the user that you associate with the key must exist on the Storwize V7000 File Module.

Updating the credentials for a storage system that is managed by a CIM agent

Change the credentials that IBM Spectrum® Control uses to authenticate to a CIM agent.

About this task

IBM Spectrum Control communicates with SMI-S providers to collect information about the following resources:

- Non-IBM® storage systems that are managed by SMI-S certified Common Information Model Object Manager (CIMOM), such as Dell EMC storage systems other than Unity, Hitachi, and NetApp
- Switches: Brocade switches that are running with a version of Fabric OS earlier than 8.2.1, monitored through Brocade Network Advisor

To update the credentials for a CIM agent for a storage system, complete the following steps. You can update the host name or IP address, user name, password, and other information. You can also update the CIM agent credentials for other devices, such as switches, from the GUI pages for those devices.

Procedure

1. In the menu bar, go to Storage > Block Storage Systems.
Information about monitored storage systems is displayed.
2. Right-click a storage system and click Connections > Modify Connection.
3. Change the CIM agent host name or IP address, the user name, or the password. Under Advanced, you can also specify the protocol, port, and namespace. Click OK.

Testing the connection to a storage system

Verify that IBM Spectrum® Control can communicate with a monitored storage system. For storage systems that are managed by a CIM agent or Storage Resource agent, the connection to the agent is tested.

About this task

To test the connection to a storage system, complete the following steps:

Procedure

1. Depending on the type of storage system that you want to test, go to Storage > Block Storage Systems, Storage > File Storage Systems, or Storage > Object Storage Systems.
Information about monitored storage systems is displayed.
2. Right-click a storage system and click Connections > Test Connection.
A message that shows the results of the test is displayed.

Collecting CIM agent logs

You can collect logs for certain IBM® CIM agents using the command line interface.

About this task

Remember: Storage systems that use the native interfaces (for example, DS8000®, the XIV®, SAN Volume Controller, and Storwize® V7000) do not use CIM agents.

Procedure

1. Change to the directory where the CIM agent is installed.
2. Run one of the following commands:

On Linux® operating systems

collectLogs.sh

On Windows operating systems

collectLogs.bat

A **collectedLogs.zip** file is created.

Important: This file is overwritten if you run the script again.

Removing storage systems

Remove storage systems that you no longer want to monitor with IBM Spectrum® Control.

About this task

To remove a storage system, complete the following steps:

Procedure

1. In the menu bar, go to Storage and select the type of storage system that you want to remove.
2. Right-click a storage system and click Remove.
3. Click Remove to confirm that you want to remove the storage system.

Hypervisors and VMware data sources

Administer the hypervisors, vCenter Server Appliance systems and vCenter Server systems that are monitored by IBM Spectrum® Control. vCenter Server Appliance systems and vCenter Server systems are data sources that can monitor multiple hypervisors. A hypervisor can be an ESX or ESXi host. Each hypervisor can host multiple virtual machines.

About this task

- [Checking permissions to browse data stores](#)
Determine if the user name that you specified for a VMware data source has permission to browse through the data stores on a hypervisor.
- [Viewing information about hypervisors](#)
View detailed information about hypervisors that are monitored by IBM Spectrum Control.
- [Updating the credentials for a hypervisor](#)
You can change the user name and password that IBM Spectrum Control uses to log in to a hypervisor. You can also change the host name or IP address of the hypervisor.
- [Removing hypervisors and VMware data sources](#)
Remove hypervisors and VMware vCenter servers that you no longer want to monitor with IBM Spectrum Control.

Checking permissions to browse data stores

Determine if the user name that you specified for a VMware data source has permission to browse through the data stores on a hypervisor.

About this task

When you add a VMware data source in IBM Spectrum® Control, the user name that you specify must have permission to browse through the data stores on VMware. IBM Spectrum Control must browse through the data stores to collect information from the hypervisors. However, the "Read Only" role as defined by VMware does not allow IBM Spectrum Control to browse the data stores. You can use the "Virtual Machine Power® User" role if you do not want to use the Administrator role, or you can create a custom role with the required permissions.

Procedure

To verify that a VMware user is assigned the correct role and privileges to monitor VMware data sources, follow these steps:

1. Ensure that the user role has the required VMware datastore permissions by completing the following steps:
 - a. Connect the vSphere Web Client to the VMware data source.
The data source can be an ESX server, a vCenter Server Appliance, or a vCenter Server.
 - b. From the Inventories view, select Hosts and Clusters.
 - c. Select a host, and click the Related Objects tab.
 - d. View the datastores by clicking the Datastores tab.
 - e. Right-click a datastore, and select File Browser. If you can view the Files tab for the datastore, your browse permission is working correctly.
2. Determine the role that is assigned to the user by logging in to the vSphere Web Client by using the administrator user ID. From the Administration view, select Roles. Verify the role name that is assigned to the user.
3. Determine the privileges that are assigned to the role by selecting the user's role and clicking Privileges. Expand the privilege groups to view the specific privileges.
4. Optional: If you must edit the privileges for the role, select the role and click the Edit role action icon. Select privilege groups or expand to select specific privileges.

What to do next

For more information about VMware user roles, go to the [VMware documentation center](#) and search for *vSphere users and permissions*.

Viewing information about hypervisors

View detailed information about hypervisors that are monitored by IBM Spectrum® Control.

Procedure

To view information about hypervisors and vCenter servers, complete the following steps:

1. Go to Servers > Hypervisors
Information about monitored hypervisors and vCenter servers is displayed.
2. Right-click a hypervisor and select View Properties to view the key properties of that hypervisor.
3. Optional: Right-click a hypervisor and select View Details to view more detailed information about that hypervisor, such as triggered alerts, data collection schedules, and information about its internal and related resources.

Updating the credentials for a hypervisor

You can change the user name and password that IBM Spectrum® Control uses to log in to a hypervisor. You can also change the host name or IP address of the hypervisor.

Procedure

1. In the menu bar, go to Servers > Hypervisors.
2. Right-click a hypervisor and select Connections > Modify Connection.
3. Update the host name or IP address, user name, or password for the hypervisor.
The user name and password must contain the following valid characters:
 - A through Z (uppercase characters)
 - a through z (lowercase characters)
 - 0 through 9 (numeric characters)
 - Special characters: ! # % & * + - / = ? ^ _ { } () . ,Restrictions:
 - User names and passwords cannot contain spaces and must have at least one character.
 - The maximum length of a user name or password is 128 characters.
 - The user name must have permission to browse the data stores on a hypervisor. For more information about permissions, see [Checking permissions to browse data stores](#).
4. Click OK.

Removing hypervisors and VMware data sources

Remove hypervisors and VMware vCenter servers that you no longer want to monitor with IBM Spectrum® Control.

Procedure

To remove hypervisors and vCenter servers, complete the following steps:

1. Go to Servers > Hypervisors
Information about monitored hypervisors and vCenter servers is displayed.
2. Right-click a hypervisor and select Remove.
The hypervisor and all its data are removed from IBM Spectrum Control immediately. Any data collection jobs and alerts are also removed.
When you remove a vCenter Server, the hypervisors that it manages are also removed from IBM Spectrum Control. However, information about the hypervisors is not removed immediately, but is retained according to the Data for missing resources setting on the History Retention page. The default setting is 14 days. If the default setting is used, all information about the hypervisors is deleted 14 days after the related vCenter Server was removed.
Tips: After a vCenter Server is removed, but before its managed hypervisors are removed according to the retention settings, the following conditions occur:
 - Any data collection jobs that are scheduled for the hypervisors fail.
 - Because data is no longer collected, any alerts that were based on that data are not generated.

Switches and fabrics

Administer the switches and fabrics that are monitored by IBM Spectrum® Control. Administering actions include adding and removing switches and fabrics, modifying connection information, and testing connections.

- [Viewing information about switches and fabrics](#)
View detailed information about switches and fabrics that are monitored by IBM Spectrum Control.
- [Updating the connection information for switches and fabrics](#)
Change the connection information that IBM Spectrum Control uses to authenticate to a data source that manages a switch or fabric.
- [Testing the connection to a switch or fabric](#)
Verify that IBM Spectrum Control can communicate with the data source that manages a switch or fabric.
- [Removing switches and fabrics](#)
Remove a switch or fabric that you no longer want to monitor with IBM Spectrum Control.

Viewing information about switches and fabrics

View detailed information about switches and fabrics that are monitored by IBM Spectrum® Control.

About this task

To view information about switches and fabrics, complete the following steps:

Procedure

1. In the menu bar in the web-based GUI, go to Network > Switches or Network > Fabrics. Information about monitored switches or fabrics is displayed.
2. Right-click a switch or fabric and click View Properties to view the key properties for the switch or fabric.

Updating the connection information for switches and fabrics

Change the connection information that IBM Spectrum® Control uses to authenticate to a data source that manages a switch or fabric.

Before you begin

IBM Spectrum Control uses REST APIs, SMI agents, and SNMP agents to collect data about switches and fabrics in your SAN. You can update the connection information for a switch and fabric to change its data source at any time. The type of data source determines the connection information that you can update.

- [Updating the connection information for a switch](#)
Change the connection information that IBM Spectrum Control uses to authenticate to a data source that manages a switch. The data source can be a REST API, SMI agent, or SNMP agent.
- [Updating the connection information for a fabric](#)
Change the connection information that IBM Spectrum Control uses to authenticate to a data source that manages a fabric. The data source can be an SMI-S provider or an SNMP agent.

Updating the connection information for a switch

Change the connection information that IBM Spectrum® Control uses to authenticate to a data source that manages a switch. The data source can be a REST API, SMI agent, or SNMP agent.

About this task

To update the connection information for a switch, complete the following steps:

Procedure

1. In the menu bar, go to Network > Switches. Information about monitored switches is displayed.
2. Right-click a switch and click Connections > Modify Connection.
Tip: If a switch is managed by multiple data sources, the menu is displayed as Connections > Modify Connection > *data_sources*. Select the data source for which you want to update the connection information. For example, a Brocade switch that is running with a version of Fabric OS earlier than 8.2.1 might have multiple SMI-S providers, and therefore multiple data sources.
3. Update the following information as required and then click OK.
The information that is displayed depends on the type of data source.

REST API

Host names or IP addresses

The IP address or host name to use to connect to the switch.

Protocol, Port

The https or http protocol and the 443 or 80 port to use to connect to the switch.

User name, Password

The user name and password for logging on to the switch.

SMI-S provider

SMI-S provider host name or IP address

The IP address or host name of the SMI-S provider that manages the switch. For Brocade switches, the SMI-S provider is on Brocade Network Advisor (BNA).

User name, Password

The user name and password for logging on to the SMI-S provider.

Advanced

Protocol, Port

The https or http protocol and the 5989 or 5988 port to use to connect to the SMI-S provider.

Namespace

The namespace that includes the class instances of the Server Profile. The interaction with the SMI-S provider when information is retrieved is determined by the namespace.

SNMPv3 agent

SNMP version

The SNMP version of the agent.

User name

The user name used to log in to the switch.
Authentication password
The password for the user logged in to the switch.
Authentication protocol
The protocol or digest used for authentication to the switch.
Encryption protocol
The protocol used for encryption.

SNMPv1 agent

SNMP version
The SNMP version of the agent.
Read community
The SNMP community string. The default is public.
Write community
The SNMP community string. The default is private.

Updating the connection information for a fabric

Change the connection information that IBM Spectrum® Control uses to authenticate to a data source that manages a fabric. The data source can be an SMI-S provider or an SNMP agent.

About this task

To update the connection information for a fabric, complete the following steps:

Procedure

1. In the menu bar, go to Network > Fabrics.
Information about monitored fabrics is displayed.
2. Right-click a fabric and click Connections > Modify Connection.
Tip: If a fabric is managed by multiple data sources, for example multiple SMI-S providers, the menu is displayed as Connections > Modify Connection > *data_sources*. Select the data source for which you want to update the connection information.
3. Update the following information as required and then click OK.
The information that is displayed depends on the type of data source.

SMI-S provider

SMI-S provider host name or IP address
The IP address or host name of the SMI-S provider that manages the switch. For Brocade switches, the SMI-S provider is on Brocade Network Advisor (BNA).
User name, Password
The user name and password for logging on to the SMI-S provider.
Advanced
Protocol, Port
The https or http protocol and the 5989 or 5988 port to use to connect to the SMI-S provider.
Namespace
The namespace that includes the class instances of the Server Profile. The interaction with the SMI-S provider when information is retrieved is determined by the namespace.

SNMPv3 agent

SNMP version
The SNMP version of the agent.
User name
The user name used to log in to the switch.
Authentication password
The password for the user logged in to the switch.
Authentication protocol
The protocol or digest used for authentication to the switch.
Encryption protocol
The protocol used for encryption.

SNMPv1 agent

SNMP version
The SNMP version of the agent.
Read community
The SNMP community string. The default is public.
Write community
The SNMP community string. The default is private.

Testing the connection to a switch or fabric

Verify that IBM Spectrum® Control can communicate with the data source that manages a switch or fabric.

About this task

To test the connection to the data source that manages a switch or fabric, complete the following steps:

Procedure

1. In the menu bar in the web-based GUI, go to **Network > Switches** or **Network > Fabrics**.
Information about monitored fabrics or switches is displayed.
2. Right-click a switch or fabric and click **Connections > Test Connection**.
A message that shows the results of the test is displayed.

Removing switches and fabrics

Remove a switch or fabric that you no longer want to monitor with IBM Spectrum® Control.

About this task

To remove a switch or fabric, complete the following steps:

Procedure

1. In the menu bar in the web-based GUI, go to **Network > Switches** or **Network > Fabrics**.
Information about monitored fabrics or switches is displayed.
2. Right-click a fabric or switch and select **Remove**.
Note: To remove a Cisco fabric, you must remove all the switches in that fabric. The Cisco fabric is then automatically removed.
3. Follow the directions that are presented in the information message.

Servers and Storage Resource agents

Administer servers and the Storage Resource agents that collect asset, status, and file system information about servers.

- **[Fixing deployments](#)**
Use the Servers page to monitor servers that are added to IBM Spectrum® Control by deploying a Storage Resource agent. You can identify agents that failed to deploy, investigate and resolve the problems that caused the deployment failure, and deploy the agents again.
- **[Canceling deployments](#)**
Use the Servers page to cancel the deployment of Storage Resource agents.
- **[Modifying deployment schedules](#)**
Use the Servers page to modify deployment schedules for Storage Resource agents.
- **[Viewing information about Storage Resource agents](#)**
View detailed information about the Storage Resource agents that are deployed on monitored resources.
- **[Viewing Storage Resource agent log files](#)**
The log files for a Storage Resource agent contain informational, warning, and error messages for the actions that were taken by the agent. You can use the content of the log files to troubleshoot any errors that might occur when a Storage Resource agent is started, processing data, or shut down.
- **[Disabling Storage Resource agents](#)**
Disable Storage Resource agents so that they no longer collect data or run IBM Spectrum Control jobs.
- **[Enabling Storage Resource agents](#)**
You can enable Storage Resource agents that are in a Disabled or Down state. After an agent is enabled, the IBM Spectrum Control server resumes communication with that agent.
- **[Testing the connection with a Storage Resource agent](#)**
Verify that the IBM Spectrum Control server can communicate with the server where a Storage Resource agent is deployed.
- **[Changing credentials for Storage Resource agents](#)**
You can change Storage Resource agent credentials, such as the user name and password that IBM Spectrum Control uses for logging on to the server where the agent is deployed.
- **[Collecting service data](#)**
Collect service data about the selected Storage Resource agent. Service data includes diagnostic information such as logs, trace files, configuration information, and computer details. Use this information to troubleshoot any errors that might occur during startup, processing, or shutdown of a Storage Resource agent.
- **[Enabling or disabling scripts for Storage Resource agents](#)**
You can enable or disable scripts that are sent from the IBM Spectrum Control server to Storage Resource agents.
- **[Enabling or disabling the monitoring of fabrics by Storage Resource agents](#)**
You can enable or disable fabric monitoring by Storage Resource agents. Fabric monitoring is enabled by default. When you enable fabric monitoring, the agent collects information about fabrics that the server is connected to.
- **[Using the help command for Storage Resource agents](#)**
The **help** command for Storage Resource agents provides information about the parameters for installing, uninstalling, and upgrading Storage Resource agents.
- **[Removing servers](#)**
You can remove servers that you no longer want to monitor with IBM Spectrum Control.
- **[Registering a Storage Resource agent with a different IBM Spectrum Control server](#)**
You can register a Storage Resource agent with a different IBM Spectrum Control server.
- **[Manually changing the Windows service logon](#)**
Change the Windows service logon for a Storage Resource agent.

- [Deployment guidelines and limitations for Storage Resource agents](#)

You must consider the following guidelines and limitations when you manage Storage Resource agents in your environment.

Related tasks

- [Upgrading Storage Resource agents](#)

Fixing deployments

Use the Servers page to monitor servers that are added to IBM Spectrum® Control by deploying a Storage Resource agent. You can identify agents that failed to deploy, investigate and resolve the problems that caused the deployment failure, and deploy the agents again.

Before you begin

To use the Fix Deployment action, you must have Administrator privileges.

About this task

When you use the Fix Deployment action, the existing agent deployment on the server is automatically overwritten when the agent is deployed again.

Use the following steps to identify and fix Storage Resource agents that failed to deploy:

1. Use the Status column on the Servers page to identify agents that failed to deploy. A status of Failed deployment indicates that an error occurred when the agent was deployed.
2. Use the deployment log to investigate the problems that prevented the agent from deploying.
Tip: The Open Logs action is not available if you select multiple server rows. The Fix Deployment action is available if you select a single server row or multiple server rows.
3. Use the Fix Deployment action to change the deployment settings for the agents and deploy the agents again.

The following examples show some of the problems that cause agent deployments to fail and the actions that you might take to resolve the problems:

Errors that do not require changes to the deployment settings

The log message indicates that the Db2® database or the Data server is not running. Start the service that is not running and use the Fix Deployment action to deploy the agent. You do not need to change the deployment settings.

Errors that require changes to the deployment settings

The log message indicates that the port number on which the agent listens for requests from IBM Spectrum Control is in use by another service. Use the Fix Deployment action to change the setting for the Port field and to deploy the agent.

Procedure

1. In the menu bar, go to Servers > Servers.
2. Locate the servers with failed agent deployments that you want to fix.
3. For each server with a status of Failed deployment, complete the following steps:
 - a. To view the error messages, right-click the server row and click Open Logs.
 - b. Investigate and resolve the errors.
4. Click a single or multiple servers with a status of Failed deployment and click Actions > Fix Deployment.
5. On the Deploy Agent page, change the settings that caused the deployment errors.
For example, if the deployment fails because there is not enough disk space at the location that is specified in the Installation path field, you might change the installation location for the agents.
If you selected multiple servers with different operating systems, separate configuration pages are displayed for agents that are deployed on Windows servers and agents that are deployed on UNIX servers.

Tip: If you select multiple servers, the following rules are used to determine the settings for the agent deployment fields:
 - a. If the servers use different authentication methods, you cannot change the authentication settings. Keep current settings is displayed in the Authentication field and the fields that are used to configure the authentication settings are hidden.
 - b. If the servers are configured with different daemon modes, you can specify the daemon mode to apply to all the selected servers.
 - c. For other fields, if the servers have the same value for the field, the value is displayed. If the servers have different values for the field, the field is blank.
6. On the Configure page, if the setting for the Location field caused a deployment error, change the field setting.
7. Schedule the deployment of the Storage Resource agents.
If you are fixing the agent deployment for multiple servers, a time span is calculated during which the agents are deployed. The agents are deployed at regular intervals during the time span to avoid excessive load on the IBM Spectrum Control server.
8. Schedule the time and frequency that probes are run for the servers.
If you are fixing the agent deployment for multiple servers, a time span is calculated during which the servers are probed.
9. Click Finish to deploy the agents.

Results

The changes are applied to the servers that have a status of Failed deployment. If you select servers that have other statuses, for example, Pending deployment, those servers are not affected by the action.

A probe is automatically run for a server after the agent is successfully deployed.

What to do next

To monitor the status of the agent deployment, check the Agent State column on the Servers page.

Canceling deployments

Use the Servers page to cancel the deployment of Storage Resource agents.

Before you begin

To use the Cancel Deployment action, you must have Administrator privileges.

About this task

Check the Agent State column on the Servers page to identify the agent deployments that you can cancel. You can cancel the agent deployment for servers with a status of Failed deployment or Pending deployment.

Procedure

1. In the menu bar, go to Servers, > Servers.
2. Locate the servers with the agent deployments that you want to cancel.
3. Click a single or multiple servers with a status of Failed deployment or Pending deployment, and then click Actions, > Cancel Deployment.

Results

The agent deployment is canceled for the servers with a status of Failed deployment or Pending deployment. If you select servers that have other statuses, for example, Deploying, those servers are not affected by the action.

When you cancel the agent deployments, the servers are removed from IBM Spectrum® Control. To add the servers again, click Deploy Agent.

Modifying deployment schedules

Use the Servers page to modify deployment schedules for Storage Resource agents.

Before you begin

To use the Modify Deployment Schedule action, you must have Administrator privileges.

About this task

Check the Agent state column on the Servers page to identify the agent deployments that you can modify. You can modify the deployment schedules for servers that have a status of Pending deployment.

The Modify Deployment Schedule action is available if you click a single server row or multiple server rows.

Procedure

1. In the menu bar, go to Servers, > Servers.
2. Locate the servers with the agent deployments that you want to modify.
3. Click a single or multiple servers with a status of Pending deployment, and then click Actions, > Modify Deployment Schedule.
4. On the Modify Deployment Schedule window, the current schedule values for the agent deployments are shown. You can change the date and time that agents are deployed.

If you are modifying the deployment schedule for multiple agents, a time span is calculated during which the agents are deployed. The agents are deployed at regular intervals during the time span to avoid excessive load on the IBM Spectrum® Control server.

Tips:

- If you select multiple servers and the servers have the same value for a field, the value is displayed. For example, if the selected servers have the same deployment date, the date is displayed. If the servers have different values for the field, the field is blank.
- The scheduled time for an agent deployment is based on the time zone of the IBM Spectrum Control server, not the time zone of the server where the agent is deployed.

5. Click Save.

Results

The deployment schedules are modified for the servers that have a status of Pending deployment. If you select servers that have a status other than Pending deployment, the changes to the deployment schedule are not applied to those servers.

Viewing information about Storage Resource agents

View detailed information about the Storage Resource agents that are deployed on monitored resources.

Procedure

To view information about a Storage Resource agent, complete the following steps:

1. In the menu bar, go to Servers, > Servers.
2. On the Servers page, right-click the server where the agent is deployed and select View Properties.
3. In the properties notebook, click the Agent tab.

Results

Detailed information about the agent is shown, such as the agent state and version, and the date and time when the agent was last updated.

If the Storage Resource agent has a state of Upgrade needed, the agent must be upgraded to the same version level as the IBM Spectrum® Control server to which it is communicating.

Viewing Storage Resource agent log files

The log files for a Storage Resource agent contain informational, warning, and error messages for the actions that were taken by the agent. You can use the content of the log files to troubleshoot any errors that might occur when a Storage Resource agent is started, processing data, or shut down.

About this task

By default, the log files are located in the following directories on the server where an agent is deployed:

Windows

C:\Program Files\IBM\TPC\agent\log\SRV1\agent.log

Linux®, UNIX, and AIX®

/opt/IBM/TPC/agent/log/computer_name/agent.log




where *computer_name* represents the name of the server where IBM Spectrum® Control is installed. If an agent communicates with more than one installation of IBM Spectrum Control, a subfolder is created for each installation. For example, if the agent communicates with servers named SRV1 and SRV2, the following folders are created:

- C:\Program Files\IBM\TPC\agent\log\SRV1\agent.log
- C:\Program Files\IBM\TPC\agent\log\SRV2\agent.log

To view the log file for a Storage Resource agent, complete the following steps:

Procedure

1. In the menu bar, go to Servers, > Servers.
2. On the Servers page, locate the server that contains the Storage Resource agent that you want to analyze.
3. Right-click the server row and select Logs, > View Agent Log.
4. Optional: To view only the log entries that have a Warning or Error status, select an option from the Show all list.
You can choose to view only entries that have the following statuses:

-  Only error entries
-  Only warning entries
-  Error and warning entries

5. Optional: To view an explanation of the message that is associated with a log entry, click the link in the ID column.

Disabling Storage Resource agents

Disable Storage Resource agents so that they no longer collect data or run IBM Spectrum® Control jobs.

About this task

You might want to disable a Storage Resource agent under the following conditions:

- The monitored server is undergoing maintenance and is unavailable. This action prevents IBM Spectrum Control from flagging the agent as "down" if it cannot reach the agent. The number of times that the server tries to contact the agent is defined by the agentErrorLimit parameter in the server.config file.
- The monitored server is busy with resource-intensive processing and you do not want to add any IBM Spectrum Control jobs to that processing load.

Procedure

To disable a Storage Resource agent, complete the following steps:

1. In the menu bar, go to Servers, > Servers.
2. On the Servers page, right-click the server where the agent is deployed and select Modify Agents, > Disable.
3. Click OK to confirm that you want to disable the agent.

The state of the agent is changed to Disabled and remains in that state until it is enabled again. You can disable agents on multiple servers at the same time.

Results

When you disable a Storage Resource agent that is deployed as a daemon service, the service is shut down, and the agent is disabled. IBM Spectrum Control no longer sends requests to the agent or contacts it for job processing. A Storage Resource agent that is deployed as a non-daemon agent runs as a stand-alone process. Because a service is not required for this type of agent, it is not necessary to shut down the agent before it is disabled.

Enabling Storage Resource agents

You can enable Storage Resource agents that are in a Disabled or Down state. After an agent is enabled, the IBM Spectrum® Control server resumes communication with that agent.

About this task

If the IBM Spectrum Control server cannot contact an agent, the agent is automatically flagged as "down". You can use the Enable action to reestablish communication between the agent and the IBM Spectrum Control server. The number of times that the IBM Spectrum Control server tries to contact the agent is specified in the agentErrorLimit parameter in the server.config file. The default value for the agentErrorLimit parameter is 3.

By default, the server.config file is located in the following directory:

Windows

C:\Program Files\IBM\TPC\Data\config

Linux® or UNIX

/opt/IBM/TPC/Data/config

Procedure

1. In the menu bar, go to Servers, > Servers.
2. On the Servers page, right-click the server where the agent is deployed and select Modify Agents, > Enable.
You can enable agents on multiple servers at the same time.
3. Click OK to confirm that you want to enable the agent.
4. If the agent is running as a daemon service, enter the user ID, password, and other credentials for the server where the agent is deployed. Click OK to start the service and enable the agent.

Results

The agent is enabled and the state of the agent is updated to reflect its current condition, such as Up or Upgrade needed. If the agent is deployed as a daemon service, the service is started when you enable the agent.

Testing the connection with a Storage Resource agent

Verify that the IBM Spectrum® Control server can communicate with the server where a Storage Resource agent is deployed.

About this task

Use the Test Connection action in the web-based GUI to verify the state of the Storage Resource agent. For example, if the agent has a state of Down or Unreachable on the Servers page, you can test the connection to verify the state of the agent.

Procedure

1. In the menu bar, go to Servers, > Servers.
2. On the Servers page, right-click the server where the Storage Resource agent is deployed and select Modify Agents, > Test Connection.
3. Optional: If the process is slow, click Close in the Testing Agent Connection window to run the operation in the background.

Results

When the operation is complete, the server status and the agent state are automatically updated on the Servers page.

Changing credentials for Storage Resource agents

You can change Storage Resource agent credentials, such as the user name and password that IBM Spectrum® Control uses for logging on to the server where the agent is deployed.

Procedure

1. In the menu bar, go to Servers, > Servers.
2. On the Servers page, right-click the server where the agent is deployed and select Modify Agents, > Update Credentials.
3. In the Enter User Credentials window, change the credentials for logging on to the server where the agent is installed.
You can change the following credentials:

User name, Password

The user name and password that IBM Spectrum Control uses for logging on to the server where the Storage Resource agent is deployed. The user name must have administrative or root privileges on the server. This action is available only for Storage Resource agents that were deployed as non-daemon services.

The user name and password must contain valid characters. You can enter the following characters:

- A - Z (uppercase characters)
- a - z (lowercase characters)
- 0 - 9 (numeric characters)
- Series of punctuation marks or special characters: ! # % & * + - / = ? ^ _ { } () . ,

Restrictions:

- User names and passwords cannot contain spaces and must have at least one character.
- The maximum length of a user name or password is 128 characters.

Certificate location

The fully qualified path of the certificate file for the Storage Resource agent, for example, `installation_dir/data/sra/operating_system/certs/sra.pem`. This file is on the computer where the IBM Spectrum Control server is installed. If the agent uses Secure Shell (SSH) protocol for communication, the certificate location field is displayed.

Passphrase

The passphrase for the certificate file. The passphrase was created when the certificate was generated.

4. Click OK to save the changes.

What to do next

Related reference

- [Protocol support for Storage Resource agents](#)

Collecting service data

Collect service data about the selected Storage Resource agent. Service data includes diagnostic information such as logs, trace files, configuration information, and computer details. Use this information to troubleshoot any errors that might occur during startup, processing, or shutdown of a Storage Resource agent.

Procedure

To collect service data for a Storage Resource agent, complete the following steps:

1. In the menu bar, go to Servers > Servers.
2. Right-click a server and select Logs > Collect Agent Logs.
A message is displayed that shows the location where the service file is stored on the IBM Spectrum® Control server.
3. In a command line or other navigation tool, go to the directory where the service file is located and unpack its contents.

Results

If the collection of service data is successful, a message is displayed that shows the location of the resulting service file (.zip). The file is stored in a directory on the same computer as the IBM Spectrum Control server. The file is in the following default directories:

- Windows operating system: `C:\Program Files\IBM\TPC\data\log\SRATraces\agent_computer_name\TPCServiceInfo.zip`
- UNIX or Linux® operating system: `/opt/IBM/TPC/data/log/SRATraces/agent_computer_name/TPCServiceInfo.zip`

Where *agent_computer_name* represents the name of the server on which a Storage Resource agent is deployed. If an agent communicates with more than one installation of IBM Spectrum Control, a subfolder is created for each installation.

If the collection of service data fails, an error message is displayed. For more information about why a data collection failed, see the server log file or the services script. These files are in the following default directories:

- Server log file (on the computer where the IBM Spectrum Control server is installed):
 - Windows operating system: `c:\Program Files\IBM\TPC\data\log`
 - UNIX or Linux operating system: `/opt/IBM/TPC/data/log`
- Services script file (on the server where the Storage Resource agent is deployed):
 - Windows operating system: `C:\Program Files\IBM\TPC\agent\service\agent_computer_name\TPCServiceInfo.html`
 - UNIX or Linux operating system: `/opt/IBM/TPC/agent/service/agent_computer_name/TPCServiceInfo.html`

Where *agent_computer_name* represents the name of the server on which the Storage Resource agent is deployed.

Enabling or disabling scripts for Storage Resource agents

You can enable or disable scripts that are sent from the IBM Spectrum® Control server to Storage Resource agents.

About this task

If you enable scripts to run, the Storage Resource agent runs the scripts that are sent from the IBM Spectrum Control server.

If you disable scripts from running, the Storage Resource agent only runs the scripts that are stored on the server where the agent is deployed. The agent does not run scripts that are sent from the IBM Spectrum Control server.

Procedure

1. In the menu bar, go to Servers, > Servers.
2. On the Servers page, right-click the server where the agent is deployed. Select Modify Agents > Enable running scripts on agent or Modify Agents > Disable running scripts on agent to enable or disable scripts from running.

Enabling or disabling the monitoring of fabrics by Storage Resource agents

You can enable or disable fabric monitoring by Storage Resource agents. Fabric monitoring is enabled by default. When you enable fabric monitoring, the agent collects information about fabrics that the server is connected to.

About this task

After you install a Storage Resource agent on a server, you can enable or disable the monitoring of fabrics that the server is connected to. If you enable fabric monitoring, the agent collects information about the SAN and zoning. If you disable fabric monitoring, the agent cannot collect fabric information or monitor fabrics that the server is connected to. If the agent is the only data source that is managing the fabric, the fabric is no longer managed. A state of Unreachable is shown for the fabric on the Fabrics page.

Procedure

1. In the menu bar, go to Servers, > Servers.
2. On the Servers page, right-click the server where the agent is deployed. Select Modify Agents > Enable Fabric Functions or Modify Agents > Disable Fabric Functions to enable or disable fabric monitoring.

Using the help command for Storage Resource agents

The **help** command for Storage Resource agents provides information about the parameters for installing, uninstalling, and upgrading Storage Resource agents.

About this task

For information about the Storage Resource agent commands, run the **help** command. Follow these steps:

1. Go to the installation location for the Storage Resource agent:

```
cd <installation_location>
```

2. Run the following command:

```
bin/Agent -help
```

3. The output from the **help** command is as follows:

```
Usage:
Agent -INSTALL
      [-COMMTYPE DAEMON -AGENTPORT portnumber]
      [-FORCE]
      -INSTALLLOC pathname
      -SERVERIP address[,address,...]
      -SERVERPORT portnumber
      [-USERID username -PASSWORD password -CERT file -PASSPHRASE phrase]

Agent -UNINSTALL
      [-FORCE]
      -SERVERNAME servername

Agent -UPGRADE
      -INSTALLLOC pathname
```

For information about how to install the Storage Resource agent using commands, see [Installing Storage Resource agents by using a command](#).

For information about how to uninstall the Storage Resource agent using commands, see [Uninstalling Storage Resource agents manually](#).

For information about how to upgrade the Storage Resource agent using commands, see [Upgrading Storage Resource agents by using a command](#).

Removing servers

You can remove servers that you no longer want to monitor with IBM Spectrum® Control.

About this task

You can use the GUI to remove servers. If a Storage Resource agent is deployed to the server, the agent is uninstalled.

When the server is removed, it is no longer monitored by IBM Spectrum Control. All the data that was collected about the server is removed from the database repository.

Tip: When you remove a server, it is only removed from IBM Spectrum Control. The server is not physically deleted from the storage environment.

Procedure

To remove a server, complete the following steps:

1. In the menu bar, go to Servers, > Servers.
2. On the Servers page, right-click the server where the agent is deployed and select Remove.
3. Click Remove to confirm that you want to remove the server.

Registering a Storage Resource agent with a different IBM Spectrum Control server

You can register a Storage Resource agent with a different IBM Spectrum® Control server.

About this task

A Storage Resource agent is registered with IBM Spectrum Control server A. You want the Storage Resource agent to point instead to IBM Spectrum Control server B.

Procedure

To register a Storage Resource agent with a different IBM Spectrum Control server, use these steps:

1. From server B, in the menu bar, go to Servers, > Servers. Click Add Server, select Deploy an agent for full server monitoring, and click Manually.
2. On the Deploy Agent page, configure deployment information for the Storage Resource agent. Specify the same port number and installation location that are used for the Storage Resource agent on server A. Select Overwrite previously installed agents.
3. On the Configure page, schedule the deployment of the Storage Resource agent and click Finish.

Results

When the deployment job completes, the Storage Resource agent is registered with server B.

What to do next

Server A can no longer communicate with the Storage Resource agent. To remove the Storage Resource agent from server A, on the Servers page in the web-based GUI, right-click the server that the Storage Resource agent is deployed on and click Remove.

Manually changing the Windows service logon

Change the Windows service logon for a Storage Resource agent.

Procedure

1. Start Windows Services.
2. On the Services window, right-click **IBM Spectrum® Control Storage Resource agent - 'C:\Program Files\IBM\TPC\'**.
3. Select Properties.
4. Click the Log On tab.
5. Change the values for This account, Password, and Confirm password to the login credentials that you want to use.
If your IBM Spectrum Control server is part of a Windows domain, change this logon to <domain>\<account>. For example: mydomain\myaccount.
Important: The Storage Resource agent requires that the domain account has local administrator privileges and the "Log on as a service" and "Act as part of the operating system" user rights.
6. Click Apply and then OK to save your changes.

Deployment guidelines and limitations for Storage Resource agents

You must consider the following guidelines and limitations when you manage Storage Resource agents in your environment.

Use the following information when you deploy Storage Resource agents:

- [Multiple Storage Resource agents that are probing or scanning the same storage resources](#)
- [Platforms that support the deployment of Storage Resource agents](#)
- [Product functions that are not available for storage devices monitored by Storage Resource agents](#)
- [Required authority for deploying Storage Resource agents](#)
- [Orphan zones](#)
- [Firewalls and Storage Resource agents deployments](#)

- [Deploying Storage Resource agents on multiple computers](#)
- [Communication between the IBM Spectrum® Control server and a Storage Resource agent](#)
- [Daemon and non-daemon services](#)
- [Port numbers for Storage Resource agents deployed as a daemon service](#)
- [Authentication between the IBM Spectrum Control server and a Storage Resource agent](#)
- [Replacing default SSL certificates](#)
- [Storage Resource agents on the same computer](#)
- [Time zones for computers monitored by Storage Resource agents](#)
- [Connections for Linux® and AIX® operating systems by using Remote Shell protocol \(RSH\)](#)
- [Deployments on Windows - NetBIOS setting](#)
- [Deployments on Windows - User Account Control \(UAC\) remote restrictions](#)

Multiple Storage Resource agents that are probing or scanning the same resources

If multiple Storage Resource agents are set up to probe or scan the same storage resources, the Storage Resource agents that was added to IBM Spectrum Control first is used for the probe or scan. Therefore, only data that is gathered by the first Storage Resource agent is shown.

Platforms that support the deployment of Storage Resource agents

For a list of platforms on which you can deploy Storage Resource agents, see the [IBM Spectrum Control interoperability matrix](#) and go to the *Agents, Servers and Browsers* section.

Product functions that are unavailable for resources that are monitored by Storage Resource agents

Before you deploy a Storage Resource agent, ensure that the product functions you want to use on the monitored resources are available for those agents. The following functions are not available for resources that are monitored by Storage Resource agents:

- Certain relational database monitoring. For list of relational databases that can be monitored by Storage Resource agents, see the [IBM Spectrum Control interoperability matrix](#) and go to the *Agents, Servers and Browsers* section.
- The reporting of HBA, fabric topology, or zoning information for fabrics that are connected to hosts that are running Linux on IBM® System z® hardware. These limitations also apply to Storage Resource agents on all guest operating systems for VMware configurations.

Required authorities for deploying and running Storage Resource agents

Before you can create deployment schedules and deploy Storage Resource agents on target computers, you must meet the following requirements:

- To create deployment schedules, you must be logged in to IBM Spectrum Control with a user ID that has the Administrator role. For information about user roles, see [Authorizing users](#).
- To deploy Storage Resource agents on target computers, you must provide a user ID that has administrative rights on those computers. You enter this ID when you create a deployment schedule. IBM Spectrum Control uses this ID to log on to the target computers and install and configure the necessary runtime files for the agents.

The user under which a Storage Resource agent (daemon or non-daemon) runs must have the following authorities on the target computers:

- On the Linux or AIX operating systems, the user must have root authority. By default, an agent runs under the user 'root'.
- On the Windows operating systems, the user must have Administrator authority and be a member of the Administrators group. By default, a Storage Resource agent runs under the 'Local System' account.

Orphan zones

Storage Resource agents do not collect information about orphan zones. An orphan zone is a zone that does not belong to at least one zoneset.

Firewalls and Storage Resource agent deployments

Before you can deploy a Storage Resource agent on a computer, you must turn off the firewall on that computer. If you do not turn off the firewall, the deployment fails.

Deploying Storage Resource agents on multiple computers

If you deploy Storage Resource agents on multiple computers at the same time, the computers must have the same administrative user ID and password. IBM Spectrum Control uses these user credentials to log on to the computers when you install Storage Resource agents.

Tip: When you deploy Storage Resource agents on multiple computers, a globally unique identifier (GUID) is created for each computer (if one does not exist).

Communication between the IBM Spectrum Control server and a Storage Resource agent

The IBM Spectrum Control server connects to a monitored computer when a Storage Resource agent is deployed and whenever a data collection schedule runs against that agent.

During deployment, the server communicates with the target computer by using one of the following protocols:

- Windows server message block protocol (SMB)
- Secure Shell protocol (SSH)
- Remote execution protocol (REXEC)
- Remote shell protocol (RSH)

After deployment, the type of communication between the server and agent on that computer depends on whether you deployed the agent as daemon service or non-daemon service.

Daemon and non-daemon services

You can deploy a Storage Resource agent as a daemon or non-daemon service:

- A Storage Resource agent that is deployed as a daemon service runs in the background on the monitored computer and listens for requests from the IBM Spectrum Control server. Connectivity between the server and agent is established by using SSL. The server and agent have their respective certificates and no additional information is required besides those certificates and the security that is provided by the SSL protocol.
- A Storage Resource agent deployed as a service on demand (non-daemon service) runs as a stand-alone executable file on the monitored computer. Communication from the server to the agent uses the same protocol that was used during the deployment of the agent. Communication from the agent to the server uses SSL.
- A Storage Resource agent that is deployed as a daemon service on AIX, Linux, and Windows servers monitors disk paths in near real-time to detect errors. When deployed as a daemon service on an AIX server, the agent also monitors disk error events in near real-time. If the Storage Resource agent detects path status changes or disk errors, they are included in the status of the disks and paths. You can define alerts so that you are notified of changes to the status of the paths on monitored disks.

Only status changes for existing paths are detected. If a new path is added, or an existing path is removed, the number of paths that is displayed is not updated immediately. The number of paths is updated after the next scheduled probe collects data.

If a disk on an AIX server has an error status and you fix the error, you might want the new status of the disk to be displayed immediately. To display the new status immediately, you must reset the status indicator for the disk. To reset the status indicator, use the **errclear** command to clear the error log. To clear

the error log, use the following syntax:

```
errclear -d H -N disk_name 0
```

For example, if you fixed an error on hdisk4, and want to display the new status immediately, run the following command:

```
errclear -d H -N hdisk4 0
```

If you do not reset the status indicator for the disk, the status changes automatically after a few hours.

For information about the **errclear** command, see [errclear Command](#).

Port numbers for Storage Resource agents deployed as a daemon service

The following port numbers are used by Storage Resource agents that are deployed as daemon service:

- 9567 (For the Storage Resource agent that is deployed on the same server as IBM Spectrum Control.)
- 9510 (For Storage Resource agents that are deployed on remote servers.)

Storage Resource agents that are deployed as a non-daemon service do not use a port.

Authentication between the IBM Spectrum Control server and a Storage Resource agent

IBM Spectrum Control requires the correct authentication information (user name, password, port, certificate location, or passphrase) for monitored computers each time it communicates with Storage Resource agents on those computers. If the authentication information changes for a host computer on which a Storage Resource agent is deployed, the authentication information for that agent must be updated by using the Modify Agents, Update Credentials action on the Servers page in the GUI.

Replacing default SSL certificates

IBM Spectrum Control provides default SSL certificates for communication between the Data server and Storage Resource agent.

IBM Spectrum Control 5.2.2 uses SSL certificates with 2048-bit encryption keys whereas previous versions of IBM Spectrum Control used 1024-bit encryption keys. If you upgrade IBM Spectrum Control from a version earlier than 5.2.2, your SSL certificates are not updated automatically. If you want to use 2048-bit encryption keys with previous versions of IBM Spectrum Control, you must replace the default SSL certificates with custom SSL certificates.

For information about how to replace SSL certificates, see [Replacing default SSL certificates for the Data server and Storage Resource agents with custom SSL certificates](#).

Storage Resource agents on the same computer

You cannot deploy a Storage Resource agent on a computer where a Storage Resource agent is already installed and pointing to the same Data server. You can deploy a Storage Resource agent on the same computer as another Storage Resource agent if those agents communicate with different Data servers and use different ports when you listen for requests.

Time zones for computers that are monitored by Storage Resource agents

The time zones of computers that are monitored by Storage Resource agents are shown as Greenwich mean time (GMT) offsets in IBM Spectrum Control reports.

For example, a computer in Los Angeles shows the following time zones in the By Computer report in Asset reporting:

```
(GMT-8:00) GMT-8:00
```

Connections for Linux and AIX operating systems by using Remote Shell protocol (RSH)

If RSH is configured to use a user ID and password, the connection fails. To successfully connect to a system by using RSH, you must set up the **.rhosts** file (in the home directory of the account). RSH must be configured to accept a login from the system that is running your application.

Deployments on Windows operating systems - NetBIOS setting

To install a Storage Resource agent on Windows targets, the Enable NetBIOS over TCP/IP option must be selected in the Control Panel settings for the computer's network connections properties. To set this option, complete the following steps:

1. Open Windows Control Panel. For information about how to open Windows Control Panel, see [Accessing administration tools](#).
2. Select Network and Dial-Up Connections, some_connection, Properties, Internet Protocol (TCP/IP), Advanced, WINS, Enable NetBIOS over TCP/IP.

To determine whether these ports are not blocked for inbound requests, see the documentation for your firewall.

To determine whether security policies are blocking the connection ports, open Administrative Tools. For information about how to open Administrative Tools, see [Accessing administration tools](#).

Depending on whether your policies are stored locally or in Active Directory, follow these directions:

Policies that are stored locally

For policies that are stored locally, complete the following steps:

1. Open Windows Administrative Services.
2. Click Local Security Policy, IP Security Policies on Local Computer.

Policies that are stored in Active Directory

For policies that are stored in Active Directory, examine the IP security policies and edit or remove filters that block the ports:

- Click Administrative Tools, Default Domain Security Settings, IP Security Policies on Active Directory.
- Click Administrative Tools, Default Domain Controller Security Settings, IP Security Policies on Active Directory.

For all Windows systems, the Server service must be running to connect to a Windows system by using the Windows protocol.

The following table lists the ports that are reserved for NetBIOS. Ensure that these ports are not blocked.

Port	Description
135	NetBIOS Remote procedure call. (Not currently used.)
137	NetBIOS name service.
138	NetBIOS datagram. (Not currently used.)
139	NetBIOS session (for file and print sharing).
445	CIFS (on Windows XP).

For Windows, shares must be shared for the Guest or Everyone accounts, and password protected sharing must be disabled. To disable password protected sharing, follow these steps:

1. Click Control Panel, Networking and Sharing Center.
2. Click Change advanced sharing settings.

3. Click the down arrow next to All Networks.
4. Select Turn off password protected sharing.
5. Click Save Changes.
6. Exit from the Control Panel.

Deployments on Windows - User Account Control (UAC) remote restrictions

To install Storage Resource agents remotely on a Windows operating system, you must disable the User Account Control (UAC) remote restrictions on the Windows operating system. User Account Control is a security component on Windows operating systems.

Tip: To disable UAC restrictions, you must modify the computer registry. Serious problems might occur if you modify the registry incorrectly. Therefore, make sure that you follow these steps carefully. For added protection, back up the registry before you modify it. Then, you can restore the registry if problems occur. For information about how to back up and restore the registry, see <http://support.microsoft.com/kb/322756/>.

To disable UAC remote restrictions, follow these steps:

1. Open the Windows Run window. For information about how to open the Run window, see [Accessing administration tools](#).
2. Enter **regedit** and click OK.
3. Locate and click the following registry subkey:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\
Policies\System
```

4. Double click the EnableLUA registry entry.
5. In the Edit DWORD (32-Bit) dialog, change the value in the Value data field from 1 to 0.
6. Click OK.
7. Exit the registry editor.

Related reference

- [Planning for Storage Resource agents](#)

SMI-S providers

Administer SMI-S providers that are associated with storage resources that are monitored by IBM Spectrum® Control. SMI-S providers enable communication between IBM Spectrum Control and certain types of storage systems and switches.

IBM Spectrum Control communicates with SMI-S providers to collect information about the following resources:

- Non-IBM® storage systems that are managed by SMI-S certified Common Information Model Object Manager (CIMOM), such as Dell EMC storage systems other than Unity, Hitachi, and NetApp
- Switches: Brocade switches that are running with a version of Fabric OS earlier than 8.2.1, monitored through Brocade Network Advisor

IBM Spectrum Control communicates directly with the following resources and does not require SMI-S providers:

- System Storage® DS8000®
- SAN Volume Controller
- The XIV®
- IBM Spectrum Accelerate
- Storwize® family of storage systems
- IBM FlashSystem® family of storage systems
- IBM Cloud Object Storage
- Dell EMC Unity
- Switches: Brocade switches that are connected to directly, that are running with Fabric OS 8.2.1 or later, and that are monitored through REST API
- Switches: Cisco switches monitored through SNMPv3 or SNMPv1

- [Verifying that an SMI agent is running](#)

You can verify that an SMI agent is running from the command line interface.

- [Replacing an SMI agent for block storage systems, fabrics, and switches](#)

You can replace the SMI agent for storage resources without interrupting the collection of performance data or losing historical data.

- [Interop namespaces for SMI-S providers for switches and storage systems](#)

This section describes the namespaces for switches and storage system SMI-S providers (also called CIM agents or CIMOMs) that are used in IBM Spectrum Control.

Related tasks

- [Replacing an SMI agent for block storage systems, fabrics, and switches](#)

Verifying that an SMI agent is running

You can verify that an SMI agent is running from the command line interface.

Procedure

Run the following command:

```
telnet <IP> <port>
```

Where <IP> is the IP address of the system where the SMI agent is installed, and <port> is the port number. By default, this is 5989 for a secure connection and 5988 for an unsecured connection.

Replacing an SMI agent for block storage systems, fabrics, and switches

You can replace the SMI agent for storage resources without interrupting the collection of performance data or losing historical data.

Before you begin

Attention: Before you begin, check whether a probe job is already in progress for the resource or is scheduled to occur while you replace the SMI agent. Plan the replacement during a time when the probe job is not occurring.

About this task

You might want to replace a SMI agent for a storage resource for the following reasons:

- The SMI agent might need to be upgraded to support data collection for the storage resource.
- The SMI agent might run on a different operating system, network, or with a different IP address.

You can replace an SMI agent on a third-party block storage system, such as a Dell EMC storage system. You cannot replace an SMI agent on IBM® System Storage® DS8000® (the DS8000 uses a native API (NAPI) connection, rather than an SMI-S provider).

To replace an SMI agent, follow these steps:

Procedure

1. Using the instructions that came with your SMI agent software, install the new SMI agent and add the storage resource that it manages. The procedure varies, depending on the type of storage resource.
2. Choose one of the following procedures:

Type of storage resource	Procedure for adding a storage resource to the SMI agent configuration
To add a storage system	See Adding storage systems
To add a fabric or a switch	See Adding fabrics and switches
3. In IBM Spectrum Control, run the Add Device Wizard for the new SMI agent.
After the wizard discovers the storage resource, a message is displayed to confirm that the data source (SMI agent) was added for monitoring. Close the window.
4. In the IBM Spectrum Control GUI, right-click the storage resource and click *Connections > Test Connection > [IP address of the new SMI agent]*.
Verify that the test was successful.
5. Remove the storage system or fabric switch from the old SMI agent.
6. Optional: Shut down the old SMI agent services or the device that runs the old SMI agent.
7. Restart the probe to continue the collection of asset, capacity, and configuration data for the resource.
8. Check the status of the probe and performance monitor to ensure that the data is being collected for the storage resource.

Results

The SMI agent is replaced. Depending on your resource or SMI agent, the old association of storage resource and SMI agent might continue to display with the new association in the Connections menu. The old SMI agent does not interfere with operation. It cannot be selected as an SMI agent for the storage resource.

What to do next

For more information, see [SMI-S providers](#) in the *IBM Spectrum Control Knowledge documentation*.

Related reference

- [SMI-S providers](#)

Interop namespaces for SMI-S providers for switches and storage systems

This section describes the namespaces for switches and storage system SMI-S providers (also called CIM agents or CIMOMs) that are used in IBM Spectrum® Control.

If you specify an incorrect namespace the following issues might occur:

- The connection test fails when the SMI-S provider is added.
- The discovery does not discover all information of the system that is managed by the SMI-S provider.
- The probe fails.
- The function that you want to perform on the system might fail (for example, collecting performance data).

For information about the interoperability namespaces for storage systems, see the [IBM Spectrum Control interoperability matrix for storage systems](#).

For information about the interoperability namespace for switches and directors, see the [IBM Spectrum Control interoperability matrix for switches](#).

SNMP agents

SNMP agents are switches and directors that communicate with IBM Spectrum® Control through SNMP. »IBM Spectrum Control supports SNMPv1, SNMPv2, and SNMPv3 for connecting to switches and sending alert notifications to SNMP trap destinations.«

About this task

IBM Spectrum Control uses SNMP to send queries across the IP network to management information bases (MIBs) supported on switches and directors. IBM Spectrum Control uses the Fibre Alliance FC Management MIB and the Fibre Channel FE MIB specifications. The queries are sent only to switches that were added to IBM Spectrum Control for use as SNMP agents. Information is collected from each switch that is configured to use SNMP. The SNMP discovery registers each switch.

You add a switch as an SNMP agent on the Network » Switches page. You can then perform actions on the switch, including:

- running an immediate probe job, or scheduling a probe job to collect data from the switch at a specified time
- viewing information about the switch, such as the condition of the switch, the status of the most recent probe job, and other information
- modifying the connection information and credentials for the switch
- removing the switch so it is no longer managed by IBM Spectrum Control.

For information about adding switches to IBM Spectrum Control, see [Adding resources](#).

- **Displaying information about an SNMP agent**
You can view information about an SNMP agent including the IP address, user name, and password.
- **Removing an SNMP agent**
To remove an SNMP agent that is being managed by IBM Spectrum Control, you must remove the switch or fabric that the SNMP agent is monitoring. Data collected by the agent is not removed from the database repository.

Displaying information about an SNMP agent

You can view information about an SNMP agent including the IP address, user name, and password.

About this task

To display information about an SNMP agent, follow this procedure:

Procedure

1. In the menu bar, go to Network » Switches. Information about monitored switches is displayed.
2. Right-click an SNMP switch and click Connections » Modify Connection.
3. The following information is displayed.

SNMPv3 agent

- SNMP version
The SNMP version of the agent.
- User name
The user name used to log in to the switch.
- Authentication password
The password for the user logged in to the switch.
- Authentication protocol
The protocol or digest used for authentication to the switch.
- Encryption protocol
The protocol used for encryption.

SNMPv1 agent

- SNMP version
The SNMP version of the agent.
- Read community
The SNMP community string. The default is public.
- Write community
The SNMP community string. The default is private.

Removing an SNMP agent

To remove an SNMP agent that is being managed by IBM Spectrum® Control, you must remove the switch or fabric that the SNMP agent is monitoring. Data collected by the agent is not removed from the database repository.

About this task

To remove an SNMP agent, follow this procedure:

Procedure

1. In the menu bar of the GUI, go to Network > Switches or Network > Fabrics.
2. Right-click a switch or fabric and select Remove.
To remove a Cisco fabric, you must remove all the switches in that fabric. The Cisco fabric is then automatically removed.
3. Follow the directions that are presented in the information message.

Starting and stopping the IBM Spectrum Control servers

You can start and stop the IBM Spectrum® Control servers in the GUI or by running scripts.

About this task

IBM Spectrum Control provides scripts for starting and stopping the servers that run within the product. To run these scripts, on a Windows operating system, you must have Administrator authority.

Tip: The default *installation_dir* is C:\Program Files\IBM\TPC.

Note: IBM Spectrum Control servers start automatically on Windows, Linux®, or AIX® operating systems when the operating system is started.

- [Starting the IBM Spectrum Control servers by using the GUI](#)
Start the IBM Spectrum Control Data server, Device server, or Alert server by using the System Management page in the GUI.
- [Starting the IBM Spectrum Control servers by using scripts](#)
Run scripts to start the IBM Spectrum Control servers on the Windows, Linux, or AIX operating systems. Note: IBM Spectrum Control servers start automatically on all platforms when the operating system is started.
- [Stopping the IBM Spectrum Control servers by using the GUI](#)
Stop the IBM Spectrum Control Data server, Device server, or Alert server by using the System Management page in the GUI.
- [Stopping the IBM Spectrum Control servers by using scripts](#)
Run scripts to stop the IBM Spectrum Control servers on the Windows, Linux, or AIX operating systems.

Related reference

- [Default locations of log files](#)


Starting the IBM Spectrum Control servers by using the GUI

Start the IBM Spectrum® Control Data server, Device server, or Alert server by using the System Management page in the GUI.

Procedure

1. In the menu bar, go to Home > System Management.
2. Click Component Servers in the Components section.
3. Click Start Server next to the server that you want to start.
Tip: To start the Web server, use scripts that are provided with the product.

Results

In the Overview section of the System Management page, the running icon  is displayed next to the server to indicate that it is running.

Related concepts

- [Checking IBM Spectrum Control status](#)

Related reference

- [Starting the IBM Spectrum Control servers by using scripts](#)
- [Stopping the IBM Spectrum Control servers by using scripts](#)

Starting the IBM Spectrum Control servers by using scripts

Run scripts to start the IBM Spectrum® Control servers on the Windows, Linux®, or AIX® operating systems. Note: IBM Spectrum Control servers start automatically on all platforms when the operating system is started.

Starting the IBM Spectrum Control servers on Windows

Important: IBM Spectrum Control provides scripts for starting and stopping the servers that run within the product. To run these scripts, you must have Administrator authority.

Tip: The default *installation_dir* is C:\Program Files\IBM\TPC.

To start the servers on the Windows operating system, enter the following commands in the following order:

Data server

installation_dir\scripts\startTPCData.bat

Device server

installation_dir\scripts\startTPCDevice.bat

Alert server

installation_dir\scripts\startTPCAAlert.bat

Export server

installation_dir\scripts\startTPCExport.bat

Web server

installation_dir\scripts\startTPCWeb.bat

Storage Resource Agent - *directory*

Tip: The Storage Resource Agent service is started on the Windows operating system by using Windows Services.

To start the Storage Resource Agent service on Windows, complete the following steps:

1. Open Windows Services. For information about how to start Services, see [Accessing administration tools](#).
2. Start the IBM Storage Resource Agent - *directory* service where *directory* is where the Storage Resource agent is installed.

Starting the IBM Spectrum Control servers on Linux or AIX

Note: The default *installation_dir* is /opt/IBM/TPC.

To start the servers on the Linux or AIX operating systems, enter the following commands in the following order:

Data server

/installation_dir/scripts/startTPCData.sh

Device server

/installation_dir/scripts/startTPCDevice.sh

Alert server

installation_dir/scripts/startTPCAAlert.sh

Export server

installation_dir/scripts/startTPCExport.sh

Web server

/installation_dir/scripts/startTPCWeb.sh

Storage Resource Agent

/installation_dir/agent/bin/agent.sh start

Related concepts

- [Checking IBM Spectrum Control status](#)

Related tasks

- [Starting the IBM Spectrum Control servers by using the GUI](#)
- [Stopping the IBM Spectrum Control servers by using the GUI](#)


Stopping the IBM Spectrum Control servers by using the GUI

Stop the IBM Spectrum® Control Data server, Device server, or Alert server by using the System Management page in the GUI.

Procedure

1. In the menu bar, go to Home > System Management.
 2. Click Component Servers in the Components section.
 3. Click Stop Server next to the server that you want to stop.
- Tip: To stop the Web server, use scripts that are provided with the product.

Results

In the Overview section of the System Management page, the error icon  is displayed next to the server to indicate that it is stopped. While a server is stopped, some product functions are not available. For example, if the Alert server is stopped, the ability to detect alert conditions on resources and send notifications is not available.

Related concepts

- [Checking IBM Spectrum Control status](#)

Related reference

- [Starting the IBM Spectrum Control servers by using scripts](#)

- [Stopping the IBM Spectrum Control servers by using scripts](#)

Stopping the IBM Spectrum Control servers by using scripts

Run scripts to stop the IBM Spectrum® Control servers on the Windows, Linux®, or AIX® operating systems.

Stopping the IBM Spectrum Control servers on Windows

Important: IBM Spectrum Control provides scripts for starting and stopping the servers that run within the product. To run these scripts, you must have Administrator authority.

Tip: The default *installation_dir* is C:\Program Files\IBM\TPC.

To stop the servers on the Windows operating system, enter the following commands in the following order:

Storage Resource agent

To stop the Storage Resource Agent service on Windows, complete the following steps:

1. Open Windows Services. For information about how to open Windows Services, see [Accessing administration tools](#).
2. Stop the IBM Storage Resource Agent - *directory* service where *directory* is where the Storage Resource agent is installed.

Web server

installation_dir\scripts\stopTPCWeb.bat

Export server

installation_dir\scripts\stopTPCExport.bat

Data server

installation_dir\scripts\stopTPCData.bat

Device server

installation_dir\scripts\stopTPCDevice.bat

Alert server

installation_dir\scripts\stopTPCAAlert.bat

Stopping the IBM Spectrum Control servers on Linux or AIX

Tip: The default *installation_dir* is /opt/IBM/TPC.

To stop the servers on Linux or AIX operating systems, enter the following commands in the following order:

Storage Resource Agent

/SRA_installation_dir/agent/bin/agent.sh stop

Web server

/installation_dir/scripts/stopTPCWeb.sh

Export server

/installation_dir/scripts/stopTPCExport.sh

Data server

/installation_dir/scripts/stopTPCData.sh.

Device server

/installation_dir/scripts/stopTPCDevice.sh

Alert server

installation_dir/scripts/stopTPCAAlert.sh

Related concepts

- [Checking IBM Spectrum Control status](#)

Related tasks

- [Starting the IBM Spectrum Control servers by using the GUI](#)
- [Stopping the IBM Spectrum Control servers by using the GUI](#)

Checking the version and license of IBM Spectrum Control

The version and license of IBM Spectrum® Control that is installed on your system determine the IBM Spectrum Control functions that are available.

About this task

Check the version of IBM Spectrum Control that is installed on your system to verify that you are using the correct level of the documentation.

Check the license that is installed on your system if documented functions are not available. The functions might be restricted to a different license.

Procedure

To check the version and license of IBM Spectrum Control that is installed, complete the following steps in the web-based GUI:

1. Click the question mark icon in the banner pane of the window to display a list of help topics.

2. From the list of help topics, select About.

Checking IBM Spectrum Control status

The System Management page shows a high-level summary of the status of the server or servers on which IBM Spectrum® Control is installed. Use the System Management page to troubleshoot problems with the system, create trace logs, and get technical support.

The following system status information is available:

- The state of component servers, such as the Data, Device, and Alert servers, and the Db2® database. In a multiple-server environment, the IBM® Cognos® Analytics 11.2.0 or later server or the Db2 database can run on a separate server from the Data, Device, and Alert servers. In such a multiple-server environment, the System Management page shows which components are installed on each server.
- A chart showing the amount of used and available file system space on the server over time. Use this chart to view the storage usage trends on the server to identify or predict performance problems. In a multiple-server environment, a separate chart is shown for the two servers.
- A set of charts showing performance information for the storage system volumes that the server writes to and reads from most often. The charts show the following information for the volumes:
 - Volume utilization
 - I/O rate
 - Data rate
 - Response time
 - Read cache hits

In a multiple-server environment, these charts are shown for each server.

- Alert conditions detected on the server or servers on which IBM Spectrum Control is installed.
- [Troubleshooting problems with the IBM Spectrum Control component and servers](#)
If IBM Spectrum Control is not running or its performance has degraded, you can use the System Management page of the IBM Spectrum Control GUI to assess the overall condition of the system. You can also view the file system capacity and volume performance trends to help you anticipate future needs and prevent problems.
- [Packaging and sending log files from the System Management page](#)
To provide trace information to IBM Support about the performance of IBM Spectrum Control, you might be asked to package and send a set of log files. One way to package and send the log files is from the Spectrum Control System Management page.

Troubleshooting problems with the IBM Spectrum Control component and servers

If IBM Spectrum® Control is not running or its performance has degraded, you can use the System Management page of the IBM Spectrum Control GUI to assess the overall condition of the system. You can also view the file system capacity and volume performance trends to help you anticipate future needs and prevent problems.

Before you begin

To view file system capacity information, the Storage Resource agent on the IBM Spectrum Control server must be running. In a multiple-server environment, a Storage Resource agent must be installed and running on the secondary server to view file system capacity information for the secondary server.

To view performance information for storage system volumes, the storage systems must be managed by IBM Spectrum Control and have performance monitors running.

About this task

The System Management page shows a high-level summary of the condition of the server or servers on which IBM Spectrum Control is installed.

Procedure

To troubleshoot problems with IBM Spectrum Control, complete the following steps:

1. In the menu bar, go to Home > System Management.
2. Use the System Management page to view the status of the IBM Spectrum Control system.
 - Check the state of each component server and the DB2® database to verify that they are all running. To examine the status and resource usage of component resources in detail and, if necessary, to restart the Data server, Device server, or Alert server, complete the following steps:
 - Click Component Servers in the Components section. View the state, memory use, and database connections for each component server.
 - Optional: If the Data server, Device server, or Alert server is not running, click the Start Server button to restart the server. If the Device server is running, but one or more of its services are not running, click the Start Services button to restart the services.
 - If the performance of the IBM Spectrum Control is slow, examine the chart for available file system space and the volume performance charts.
 - Check whether there are any alerts for the server or servers on which IBM Spectrum Control is installed. The Alerts link in the Overview section shows the number of alerts and the greatest alert severity. Click Alerts in the Overview section to view the alerts.
3. Optional: You can also view the status of the product servers on Windows:
 - a. On the Windows desktop, click Start > Control Panel > Administrative Tools > Services.
Tip: For information about how to view information about services on different versions of Windows, see [Accessing administration tools](#).
 - b. On the Services window, locate the names of the server services. For example, the service for the Alert server is IBM Spectrum Control - Alert Server.
 - c. View the Status column to determine if the service is running or stopped.
 - d. Optional: If a server is not running and you want to restart it, right-click the service name for that server and click Start.

Packaging and sending log files from the System Management page

To provide trace information to IBM® Support about the performance of IBM Spectrum® Control, you might be asked to package and send a set of log files. One way to package and send the log files is from the Spectrum Control System Management page.

Before you begin

To package and send IBM Spectrum Control log files, you must be assigned to the Administrator role and open a support ticket with IBM Support. You can open a support ticket in IBM Spectrum Control as described in the following procedure or on the IBM Support website. For information about how to open a support ticket on the IBM Support website, go to [Getting support](#).

Tip: Before you package the log files, you can, optionally, adjust the level of trace recording for each component server. By selectively setting trace levels, you can provide IBM Support with more information on particular component servers that are the suspected source of the problem. You can also reduce the trace level for a particular component server to improve system performance.

About this task

When the IBM Spectrum Control component servers are running, they write trace information to log files. The log files contain trace information for component servers such as the Data server, Device server, Web server, and Alert server.

On the System Management page, you can collect the log files from all the component servers and package them into a single compressed file. You can then send the file to IBM Support automatically using FTP or manually after packaging is complete.

You can also package the log files from the command line. For more information, go to [Packaging log files from the command line and sending them to IBM Support](#).

Tip: Only one version of the log file package is retained at a time. When you create a new package, the previous package is overwritten.

Procedure

To package log files on the System Management page and send them to IBM Support, complete the following steps:

1. In the menu bar, go to Home, > System Management.
2. Optional: To adjust the trace level for any of the component servers, complete the following steps:
 - a. Click Component Servers in the Components section. In the Component Servers pane, you can view performance information for each component server. The current trace level for each of the services on a component server is displayed.
 - b. Adjust the trace level for a component server by clicking Off or On for each of the component's services.

Tip: If you adjust the trace level for one or more component servers, wait at least several hours before you package the log files. If you package the log files too soon after you adjust the trace levels, the logs will not contain enough traces at the new level. Before you create the logs, you might want to try to re-create the problem so that the log files contain the relevant trace data.
3. Click Get Support in the General section.
4. If you do not already have a support ticket associated with your issue, click **Open Support Ticket** on the Get Support page to open a support ticket. A support ticket is used by IBM Support to track and manage your issue and is required when you upload a log file package.
5. Click Collect Log.
6. Optional: If you wish to automatically send your logs using FTP, under Option 1: Automatic Upload, enter your support ticket number and your email address and then click Collect and Upload log.
7. Optional: If you prefer not to automatically send your logs using FTP, click Collect and Download Log under Option 2: Manual Upload.
8. Click Close to close out of the Creating Logs and Collect Log windows.
9. If you selected Collect and Download Log in step 7, after the log package is created, select the appropriate manual upload method under Manual Upload and complete the indicated steps.

Note: The manual upload methods are also detailed on [Getting support](#).

For more information about resolving any FTP-related problems when you upload the package, see [Troubleshooting FTP transfers](#).

Results

It can take 20 minutes or longer for IBM Spectrum Control to generate and package the log files. You can do other work or log out of IBM Spectrum Control while the package is being created. When the process completes, you can download the package by clicking the provided link on the System Management page.

Depending on the environment, the size of the log file package can vary. Its size is determined by the following factors:

- How frequently the product is used
- The number of resources that are monitored, the type of data that is being collected, and how frequently that data is collected
- The length of time that the product has been up and running

For example, if the product monitors five storage systems over a period of three days, and collects asset and performance each day, the size of the package might be 200 - 300 MB.

Related tasks

- [Troubleshooting FTP transfers](#)
- [Restoring the database](#)

Related reference

- [Getting support](#)
- [Backing up the database](#)

Increasing the memory allocation for the Data server

If the data memory that is allocated for your Data server is insufficient, you can increase the memory. The default maximum memory for the Data server is set to 1024 MB.

About this task

You cannot increase the memory for the Device server. The memory for the Device server is set to the maximum heap size for the JVM.

- [Increasing the memory allocation for the Data server that is running on AIX](#)
Increase the memory allocation for the Data server that is running on AIX®.
- [Increasing the memory allocation for the Data server that is running on Linux](#)
Increase the memory allocation for the Data server that is running on Linux®.
- [Increasing memory allocation for Data server that is running on Windows](#)
Increase the memory allocation for the Data server that is running on Windows.

Increasing the memory allocation for the Data server that is running on AIX

Increase the memory allocation for the Data server that is running on AIX®.

About this task

To increase the memory that is allocated for the Data server, complete the following steps:

Procedure

1. Log on as a user with root authority.
2. Stop the Data server. From the command line, run the following command:

```
/TPC_install_directory/scripts/stopTPCData.sh
```

Where *TPC_install_directory* is the installation directory. The default directory is */opt/IBM/TPC*.

3. Using a text editor, open the */TPC_install_directory/data/server/tpcdsrv1* file.
4. Locate the following line:

```
exec $JAVAEXE -Dsun.net.inetaddr.ttl=300 -Xrs -XmxXXXXm  
-cp $CLASSPATH com.tivoli.itsrm.server.Server &
```

where XXXX is the memory allocated for the Data server. The default is **1024m** (1024 MB).

5. Increase the memory that is allocated for the Data server. For example, to increase the memory to 1536 MB, change the line to read as follows:

```
exec $JAVAEXE -Dsun.net.inetaddr.ttl=300 -Xrs -Xmx1536m  
-cp $CLASSPATH com.tivoli.itsrm.server.Server &
```

6. Save the modified *tpcdsrv1* file.
7. Start the Data server by running the following command:

```
/TPC_install_directory/scripts/startTPCData.sh
```

Increasing the memory allocation for the Data server that is running on Linux

Increase the memory allocation for the Data server that is running on Linux®.

Procedure

To increase the memory that is allocated for the Data server, complete the following steps:

1. Log on as a user with root authority.
2. Stop the Data server.
3. From the command line, run the following command:

```
/installation_dir/scripts/stopTPCData.sh
```

Where *installation_dir* is the installation directory. The default directory is */opt/IBM/TPC*.

4. Using a text editor, open the */installation_dir/data/server/tpcdsrv1* file.
5. Locate the following line:

```
exec $JAVAEXE -Dsun.net.inetaddr.ttl=300 -Xrs -XmxXXXXm  
-cp $CLASSPATH com.tivoli.itsrm.server.Server &
```

Where XXXX is the memory that is allocated for the Data server. The default is **1024m** (1024 MB).

6. Increase the memory that is allocated for the Data server. For example, to increase the memory to 1536 MB, change the line to read as follows:

```
exec $JAVAEXE -Dsun.net.inetaddr.ttl=300 -Xrs -Xmx1536m  
-cp $CLASSPATH com.tivoli.itsrm.server.Server &
```

7. Save the modified *tpcdsrv1* file.
8. Start the Data server by running the following command:

```
/installation_dir/scripts/startTPCData.sh
```

Increasing memory allocation for Data server that is running on Windows

Increase the memory allocation for the Data server that is running on Windows.

About this task

To increase the memory that is allocated for the Data server, complete the following steps:

Procedure

1. Open the Run window. For information about how to open the Run window, see [Accessing administration tools](#).
2. Type `regedit` and click OK. The Registry Editor window is displayed.
3. Expand the HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\IBM\TSRM\1 in the Registry Editor window.
4. Right-click the 1 folder and click New\String Value.
5. Type `SRVJPARMS` as the name of the string.
6. Right-click the name of the string and click Modify.
7. Enter `-Xmx.XXXM` in the Value data field, where XXXX represents the number of megabytes for the server maximum heap size. Click OK.
The default size is 1024 MB. The largest possible value for the maximum heap size is 1536 MB. If the value is set to something larger than 1536, that value is ignored and 1536 MB is used as the maximum heap size.
8. Stop and restart the server to have the changes take effect. To stop the server, complete the following steps:
 - a. Open Windows Services. For information about how to open Services, see [Accessing administration tools](#).
 - b. Right-click IBM Spectrum® Control - Data Server and click Stop.
To restart the server, right-click the service and click Start.

Changing passwords

IBM Spectrum® Control provides a GUI and non-GUI password tool; however, both tools achieve the same purpose.

Attention: For Linux® and AIX® operating systems, it is recommended that you use the GUI password tool. However, if you use the non-GUI password tool for future updates, then *do not* use the GUI password tool to update your passwords on the same system at a later time or the IBM Spectrum Control servers might not operate properly.

If you installed IBM Spectrum Control and used the same Db2® user ID and password for the items IBM Spectrum Control requires, then when you change the Db2 password, you must also change the passwords for the items that the Db2 password applies to.

The Db2 administrative password might also apply to the following items:

- The database administration user ID and password (for the Data or Device server to connect to the database).
- The database user ID and password to create the database schema.
- The host authentication password (for the Storage Resource agents to communicate with the Device server).
- The Storage Resource agent service login user ID and password (for Windows only, if this user ID does not exist).
- [Changing passwords by using the password tool](#)
Use the password tool to change the passwords for Db2 and IBM Spectrum Control so that they can continue to authenticate to one another whenever you change a password.
- [Changing passwords on AIX and Linux systems using the Command Line Interface \(CLI\)](#)
Use the `changepasswords_noX.sh` script to change the passwords for IBM Spectrum Control on an AIX or Linux server that does not have the X Windows System installed.
- [Changing passwords on Windows systems from the Command Line Interface \(CLI\)](#)
Use the `changepasswords.bat` script to change the stored passwords of the user IDs used by IBM Spectrum Control.

Changing passwords by using the password tool

Use the password tool to change the passwords for Db2® and IBM Spectrum® Control so that they can continue to authenticate to one another whenever you change a password.

Attention: For Linux® and AIX® operating systems, it is recommended that you use the GUI password tool. However, if you use the non-GUI password tool for future updates, then *do not* use the GUI password tool to update your passwords on the same system at a later time or the IBM Spectrum Control servers might not operate properly.

If you are logged on to IBM Spectrum Control by using a domain user account, which is also a member of the local administrator group, when you run the change password tool, passwords are not updated. If you run the tool by using a local OS user account, and an error occurs, complete these steps.

To run the password tool when you log in by using a domain user account, choose one of the following methods:

- Right-click the `changepasswords.bat` file and select Run as administrator.
Or
 1. Click Start\All Programs\Open Administrative Tools\Local Security Policy. For information about how to open Administrative Tools, see [Accessing administration tools](#).
 2. On the Local Security Policy window, disable User Account Control: Run all administrators in Admin Approval Mode.
 3. Restart your computer.
- [Single server installation where components use the same logon credentials](#)
Use the password tool to change the password for IBM Spectrum Control when it is installed on a single server and the Common User and the Db2 User are

identical and use the same logon credentials. The credentials are usually *db2admin* on a Windows operating system or *db2inst1* on AIX and Linux operating systems.

- [Single-server installation where components use different logon credentials](#)

Use the password tool to change the passwords for IBM Spectrum Control when it is installed on a single server and the Common User and the Db2 User are different and use different logon credentials.

- [Multiple-server installation where Db2 is remote](#)

Use the password tool to change the passwords for IBM Spectrum Control when the IBM Spectrum Control database repository and the IBM Spectrum Control servers are installed on different servers.

Single server installation where components use the same logon credentials

Use the password tool to change the password for IBM Spectrum® Control when it is installed on a single server and the Common User and the Db2® User are identical and use the same logon credentials. The credentials are usually *db2admin* on Windows operating system or *db2inst1* on AIX® and Linux® operating systems.

Before you begin

Before you use the password tool, ensure that you know the existing password or passwords that you want to change. Stop the IBM Spectrum Control servers.

Use the following steps to change the Db2 password in the Windows, AIX, or Linux operating systems. Then, use the IBM Spectrum Control password tool to update the IBM Spectrum Control servers to use the new Db2 password.

Procedure

To change a password in the Windows operating system, follow these steps:

1. Open the Control Panel. For more information, see [Accessing administration tools](#).

- a. Select User Accounts and click Change account type.
- b. Select the user account of the password that you want to change.
- c. Select Change password.
- d. Enter and confirm the new password and click Change password.

If the password, for the Windows domain user ID that you used as the IBM Spectrum Control Common User expired, change that password in the Windows domain before you continue. If you see the following error or a similar error after you run the password tool, verify that the password for your Common User Windows domain user ID did not expire:

```
com.tivoli.itsrm.tools.changepasswords.ChangePasswords error  
SEVERE: The DB2 password is invalid.
```

If necessary, change the Common User Windows domain password and run the password tool again.

To change a password in the AIX or Linux operating system, follow these steps:

2. Log in as the root user.

- a. Run the following command:

```
passwd username
```

Where *username* is the user whose password you want to change.

- b. Enter the new password and confirm that new password is correct.

To use the IBM Spectrum Control password tool to update the IBM Spectrum Control servers to use the new password, follow these steps:

3. Open a command prompt and change the directory to the following directory:

Windows operating systems

```
Spectrum_Control_Installation_dir\service
```

Linux or UNIX operating systems

```
Spectrum_Control_Installation_dir/service
```

4. Start the password tool by running the following command:

Windows operating systems

```
changepasswords.bat
```

Linux or UNIX operating systems

```
./changepasswords
```

5. Select Change the IBM Spectrum Control and DB2 Passwords and click OK..

6. Enter and then confirm the same new password that you entered when you changed the password in the operating system in Step 1 or Step 2. Verify that the Restart servers option is selected.

7. Click OK.

8. In the Confirm password change window, click Yes.

9. When the tool finishes, click Back To Main.

10. Click Exit program..

Tip: To verify that the password changes were successful, review the log file that is located in the *Spectrum_Control_installation_dir\service\log* directory.

Related tasks

- [Single-server installation where components use different logon credentials](#)
- [Multiple-server installation where Db2 is remote](#)

Related reference

- [Stopping the IBM Spectrum Control servers by using scripts](#)
- [Changing passwords by using the password tool](#)

Single-server installation where components use different logon credentials

Use the password tool to change the passwords for IBM Spectrum® Control when it is installed on a single server and the Common User and the Db2® User are different and use different logon credentials.

Before you begin

Before you use the password tool, ensure that you know the existing password or passwords that you want to change. Stop the IBM Spectrum Control servers.

Use the following steps to change the passwords in the Windows, AIX®, or Linux® operating systems. Then, use the IBM Spectrum Control password tool to update the IBM Spectrum Control servers to use the new passwords.

Procedure

To change a password in the Windows operating system, follow these steps:

1. Open the Control Panel. For more information, see [Accessing administration tools](#).
 - a. Select User Accounts and click Change account type.
 - b. Select the user account of the password that you want to change.
 - c. Select Change password.
 - d. Enter and confirm the new password and click Change password.

If the password, for the Windows domain user ID that you used as the IBM Spectrum Control Common User expired, change that password in the Windows domain before you continue. If you see the following error or a similar error after you run the password tool, verify that the password for your Common User Windows domain user ID did not expire:

```
com.tivoli.itsrm.tools.changepasswords.ChangePasswords error  
SEVERE: The DB2 password is invalid.
```

If necessary, change the Common User Windows domain password and run the password tool again.

To change a password in the AIX or Linux operating system, follow these steps:

2. Log in as the root user.
 - a. Run the following command:

```
passwd username
```

Where *username* is the user whose password you want to change.

- b. Enter the new password and confirm that new password is correct.

To use the IBM Spectrum Control password tool to update the IBM Spectrum Control servers to use the new passwords, follow these steps:

3. Open a command prompt and change the directory to the following directory:

Windows operating systems

```
Spectrum_Control_Installation_dir\service
```

Linux or UNIX operating systems

```
Spectrum_Control_Installation_dir/service
```

4. Start the password tool by running the following command:

Windows operating systems

```
changepasswords.bat
```

Linux or UNIX operating systems

```
./changepasswords
```

5. To change the Common User's password in IBM Spectrum Control, do the following in the password tool:
 - a. Select Change IBM Spectrum Control Passwords and click OK.
 - b. Enter and then confirm the same new password that you entered when you changed the Common User's password in the operating system in Step 1 or Step 2.
 - c. Click OK.
 - d. In the Confirm password change window, click Yes.
 - e. When the process is completed, click Back to Main.
6. To change the Db2 User's password in IBM Spectrum Control, do the following in the password tool:

- a. Select Change DB2 password and click OK.
 - b. Enter and confirm the same new password as you entered when changing the Db2 User's password in the operating system in Step 1 or Step 2. Verify that the Restart servers option is selected.
 - c. Click OK.
 - d. In the Confirm password change window, click Yes.
 - e. When the process is completed, click Back to Main.
7. Click Exit Program.

Tip: To verify that the password changes were successful, review the log file that is located in the *Spectrum_Control_installation_dir*\service\log directory.

Related tasks

- [Single server installation where components use the same logon credentials](#)
- [Multiple-server installation where Db2 is remote](#)

Related reference

- [Stopping the IBM Spectrum Control servers by using scripts](#)
- [Changing passwords on Windows systems from the Command Line Interface \(CLI\)](#)

Multiple-server installation where Db2 is remote

Use the password tool to change the passwords for IBM Spectrum® Control when the IBM Spectrum Control database repository and the IBM Spectrum Control servers are installed on different servers.

Before you begin

For this procedure, the terms *Server A* and *Server B* denote the two servers. *Server A* has Db2® and the IBM Spectrum Control database repository installed. *Server B* has the IBM Spectrum Control servers installed.

Before you use the password tool, ensure that you know the existing passwords that you want to change. Stop the IBM Spectrum Control servers on *Server B*.

Use the following steps to change the passwords in the Windows, Linux®, or AIX® operating system on *Server A* and *Server B*. Then, use the IBM Spectrum Control password tool on *Server A* and *Server B* to update IBM Spectrum Control to use the new passwords.

Procedure

To change a password in the Windows operating system, follow these steps:

1. Open the Control Panel.

For more information, see [Accessing administration tools](#).

- a. Select User Accounts and click Change account type.
- b. Select the user account of the password that you want to change.
- c. Select Change password.
- d. Enter and confirm the new password and click Change password.

If the password for the Windows domain user ID that you used as the IBM Spectrum Control Common User expired, change that password in the Windows domain before you continue. If you see the following error or a similar error after you run the password tool, verify that the password for your Common User Windows domain user ID did not expire:

```
com.tivoli.itsrm.tools.changepasswords.ChangePasswords error
SEVERE: The DB2 password is invalid.
```

If necessary, change the Common User Windows domain password and run the password tool again.

To change a password in the AIX or Linux operating system, follow these steps:

2. Log in as the root user.

- a. Run the following command:

```
passwd username
```

Where *username* is the user whose password you want to change.

- b. Enter the new password and confirm that new password is correct.

To use the IBM Spectrum Control password tool to update IBM Spectrum Control to use the new passwords, follow these steps:

3. On *Server A*, open a command prompt and change the directory to the following directory:

Windows operating systems

```
Spectrum_Control_Installation_dir\service
```

Linux or UNIX operating systems

```
Spectrum_Control_Installation_dir/service
```

4. On *Server A* start the password tool by running the following command:

- For Windows operating system:

```
changepasswords.bat
```

- For Linux or AIX operating system:

changepasswords

5. Select Change DB2 password and click OK.
6. Enter and confirm the same new password as you entered when changing the Db2 User's password in the operating system in Step 1 or Step 2. Verify that the Restart servers option is selected.
7. Click OK.
8. In the Confirm password change window, click Yes.
9. When the process is completed, click Back to Main.
10. Click Exit Program.
11. On *Server B*, open a command prompt window and change the directory to the following directory:

Windows operating systems

Spectrum_Control_Installation_dir\service

Linux or UNIX operating systems

Spectrum_Control_Installation_dir/service

12. On *Server B* start the password tool by running the following command:

- For Windows operating system:

changepasswords.bat

- For Linux or AIX operating system:

changepasswords

13. To change the Common User's password in IBM Spectrum Control on *Server B*, do the following in the password tool:
 - a. Select Change IBM Spectrum Control Passwords and click OK.
 - b. Enter and then confirm the same new password that you entered when you changed the Common User's password in the operating system in Step 1 or Step 2.
 - c. Click OK.
 - d. In the Confirm password change window, click Yes.
 - e. When the process is completed, click Back to Main.
14. To change the Db2 User's password in IBM Spectrum Control on *Server B*, do the following in the password tool:
 - a. Select Change DB2 password and click OK.
 - b. Enter and confirm the same new password as you entered when you changed the Db2 User's password in the operating system on *Server A* in Step 1 or Step 2. Verify that the Restart servers option is selected.
 - c. Click OK.
 - d. In the Confirm password change window, click Yes.
 - e. When the process is completed, click Back to Main.
15. Click Exit Program.

Tip: To verify that the password changes were successful, review the log file that is located in the *Spectrum_Control_installation_dir\service\log* directory.

Related tasks

- [Single server installation where components use the same logon credentials](#)
- [Single-server installation where components use different logon credentials](#)

Related reference

- [Stopping the IBM Spectrum Control servers by using scripts](#)
- [Changing passwords on Windows systems from the Command Line Interface \(CLI\)](#)

Changing passwords on AIX and Linux systems using the Command Line Interface (CLI)

Use the `changepasswords_noX.sh` script to change the passwords for IBM Spectrum® Control on an AIX® or Linux® server that does not have the X Windows System installed.

Before you begin

Attention: It is recommended that you use the GUI password tool. However, if you use the non-GUI password tool for future updates, then *do not* use the GUI password tool to update your passwords on the same system at a later time or the IBM Spectrum Control servers might not operate properly. Before you use the `changepasswords_noX.sh` script, ensure that you know the existing passwords that you want to change. Stop the IBM Spectrum Control servers.

Use the following steps to change the Db2® password in the Windows, AIX, or Linux operating systems. Then, use the IBM Spectrum Control `changepasswords_noX.sh` script to update the IBM Spectrum Control servers to use the new passwords.

Procedure

To change a password in the AIX or Linux operating system, follow these steps:

1. Log in as the root user.
 - a. Run the following command:

```
passwd username
```

Where *username* is the user whose password you want to change.

- b. Enter the new password and confirm that new password is correct.

To use the IBM Spectrum Control `changepasswords_noX.sh` script to update the IBM Spectrum Control servers to use the new passwords, follow these steps:

2. Open a command prompt and change the directory to the following directory:

```
Spectrum_Control_Installation_dir/service
```

3. Start the `changepasswords_noX.sh` script by running the following command:

```
./changepasswords_noX.sh
```

4. Enter the corresponding option for the password that you want to change: `tpc` or `db`.
5. Enter the same new password that you entered when you changed the user password in the operating system in Step 1.
6. When you are finished using the `changepasswords_noX.sh` script, enter `quit` in the prompt window.

Changing passwords on Windows systems from the Command Line Interface (CLI)

Use the `changepasswords.bat` script to change the stored passwords of the user IDs used by IBM Spectrum® Control.

Before you use the `changepasswords.bat` to update the passwords in IBM Spectrum Control, you must change the passwords for the user IDs on the Windows operating system.

Note: The `changepasswords.bat` file is located in the `installation_dir/service` directory.

```
➔ changepasswords.bat -sc -tpc -db -p password -i -r
```

Required Parameters:

- sc OR -tpc
Changes the password stored for the user ID for IBM Spectrum Control.
- db
Changes the password stored for the database administrator user ID.
- p
Enter the new password.
- i
The new password is read from standard input, that is, you enter the password when prompted by the `changepasswords.bat` script.

Optional Parameters:

- r
Restart the IBM Spectrum Control server services.

Granting local administrative privileges to a domain account

Automatically grant administrative privileges to Windows domain accounts. The user account for the Storage Resource agent requires local administrative rights. Because these rights are not necessarily guaranteed for domain users in a Windows domain environment, you are shown how to grant local administrative rights to domain users. Using this procedure, you do not have to manually process each machine in the domain.

About this task

Note: These steps are for a Windows system that is a member of a Windows domain and not for the Windows Domain Primary Domain Controller. To use Group Policy to grant local administrative privileges to a domain account, complete the following steps:

Procedure

1. On the domain controller, go to Administrative Tools > Active Directory Users and Computers (you must be running with Domain Administrator privileges).
2. Right-click on the Organizational Unit (OU) upon which you want to apply the Group Policy. Click Properties.
3. The Group Policy Properties panel is displayed. Select the Group Policy tab and click New to create a Group Policy.
4. Designate a name for the new Group Policy. Select the new Group Policy and click Edit.
5. The Group Policy Object Editor panel is displayed. Go to New Group Policy Object *your_policy* > Computer Configuration > Windows Settings > Security Settings > Restricted Groups. Right-click Restricted Groups, and then click Add Group.
6. For example, name the new group `Administrators`. Under Properties, add the user `Administrator`, and the domain accounts or groups upon which you want the Group Policy in effect for. For example, you can add `TPC\storageadmin`, `TPC\storagegroup`, and `TPC\TestGroup`. Click OK.
7. Add these user rights to the domain account:
 - Act as part of the operating system
 - Log on as a serviceIn the Group Policy Object Editor, go to New Group Policy Object *your_policy* > Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignments. In the content pane, select "Log on as a service" and double-click. Add the domain user for whom you are granting user rights and click

OK. Repeat this step for "Act as part of the operating system."

8. The group policy is now enforced for the Organizational Unit to include the domain accounts and groups specified under the local Administrators group on each computer in the Organizational Unit. In addition, the domain user has been granted the necessary rights. To verify the user rights, log in to a domain computer and open the Computer Management console. Select Groups, double-click the Administrators group, and verify the membership of the domain users.

Collecting diagnostic information about IBM Spectrum Control

You can use the service tool to collect diagnostic information about IBM Spectrum® Control. The tool detects the system configuration, collects the applicable information, and creates a compressed file that can be sent to IBM® Software Support.

- [Service tool overview](#)
The service tool collects information from all installed IBM Spectrum Control components. The tool detects the system configuration, collects the applicable information, and creates a compressed file that can be sent to IBM Software Support.
- [Packaging log files from the command line and sending them to IBM Support](#)
To provide trace information to IBM Support about the performance of IBM Spectrum Control, you might be asked to package and send a set of log files. You can run the service tool from the command line to create the log package and then manually upload it to IBM Support.
- [Creating a compressed file for a Storage Resource agent](#)
Run the service tool on Storage Resource agents that were deployed by using the web-based GUI to create a compressed file that can be sent to IBM Software Support.
- [How to customize the service tool](#)
The default behavior of the service tool is to collect data about all IBM Spectrum Control components, but you can use the service tool to collect data about specific components. You can also use command-line parameters to specify a location to place the data that is collected, specify that the data is compressed, or to specify both.

Related tasks

- [Creating a compressed file for a Storage Resource agent](#)
- [Packaging log files from the command line and sending them to IBM Support](#)

Related reference

- [Service tool overview](#)
- [How to customize the service tool](#)

Service tool overview

The service tool collects information from all installed IBM Spectrum® Control components. The tool detects the system configuration, collects the applicable information, and creates a compressed file that can be sent to IBM® Software Support.

The service tool collects the following information:

- Host name
- IP address and configuration information
- Operating system and version. On the Windows operating system, a `msInfo.txt` report is also generated
- Java™ home, version, and class path
- Java Virtual Machine (JVM) implementation name and version
- Protocol statistics
- Internet Protocol network connections for IBM Spectrum Control, including listening ports
- Diagnostic information about the system and its services
- Listing of all library files, for example, server and library and agent and library
- HOSTS file
- IBM Spectrum Control version and license files

When the service tool is run on the system where the Data server or the Device server are installed, it also collects the following information:

- For the Data server, information about all of the remote and local graphical user interfaces (GUIs) that are associated with it
- For the Export server, information about this component configuration and logs
- For the Device server, Alert server, and web server, information about their profiles in IBM WebSphere® Application Server Liberty
- For the Data server, Device server, Alert server, web server, and data collector - all JavaCore files, the most recent Java heap dump file, and the most recent snap trace file.

The service tool can collect Java core dump files for these IBM Spectrum Control components, but it *does not* by default. The values for the `javaCoreFiles`, `heapDumpFiles`, `snapFiles`, and `coreFiles` parameters in the `installation_dir/service/service.properties` file dictate which types of files and how many of each file type are collected by the service tool.

- All applied interim fixes
- Installation logs
- The contents of the `log` and `logs` directory, including subdirectories
- The contents of the `conf` and `config` directory
- Directory listing of the `lib` and `bin` directory
- The contents of the `log` and `conf` subdirectories of the `web` directory
- For the IBM Spectrum Control GUI, information about its profile in the embedded WebSphere Application Server Liberty
- Information from the `ipconfig /all` command on Windows operating systems
- Information from the `ipconfig -a` command on Linux® and AIX® operating systems
- Information from the `netstat -an` command on all operating systems

When the service tool is run on the system where the database repository is installed, it also collects the DB2® support information. When the service tool runs on the Storage Resource agent computer, it collects the following information:

- All applied interim fixes
- Everything in the **config**, **log**, **nls**, **output**, and **service** directories, including subdirectories
- Everything in the **opt/IBM/CAP** directory on Linux and AIX operating systems
- Directory listing of the **ProgramData\Application Data\IBM\CAP** directory on Windows operating systems
- Directory listing of the **agent** directory
- Directory listing of the **bin** directory
- Listing of version numbers for the Storage Resource agent component
- Information from the **ipconfig /all** command on Windows operating systems
- Information from the **ifconfig -a** command on Linux and AIX operating systems
- Information from the **netstat -an** command on all operating systems

By default, the service data is collected in one of the following directories:

For Windows operating systems:

installation_dir\service\data

For Linux and AIX operating systems:

installation_dir/service/data

For more information about changing the default directory, see [How to customize the service tool](#).

You can run the service tool on IBM Spectrum Control regardless of whether you configured it on a single server or on multiple servers. The service tool automatically recognizes the installed components and collects service data about them. For more information about running the service tool for servers, see [Packaging log files from the command line and sending them to IBM Support](#).

Packaging log files from the command line and sending them to IBM Support

To provide trace information to IBM® Support about the performance of IBM Spectrum® Control, you might be asked to package and send a set of log files. You can run the service tool from the command line to create the log package and then manually upload it to IBM Support.

Before you begin

To package IBM Spectrum Control log files, you must have administrator authority on Windows operating systems or root authority on AIX® and Linux® operating systems. Before packaging or sending the log files, you must first open a support ticket with IBM Support. For information, see [Getting support](#).

In order to run the service tool when you log in by using a Windows domain user account, you must grant Db2® SYSADM authority to that Windows domain user account.

To run the service tool when you log in by using a domain user account, choose one of the following methods:

- Right-click the **service.bat** file and select Run as administrator.
Or
- Click Start > All Programs > Open Administrative Tools > Local Security Policy. For information about how to open Administrative Tools, see [Accessing administration tools](#).
- On the Local Security Policy window, disable User Account Control: Run all administrators in Admin Approval Mode.
- Restart your computer.

About this task

The compressed file contains data about the following IBM Spectrum Control components: Alert server, Data server, Device server, Web server, Storage Resource agent, DB2®, CLI, and installation.

Tip: You can customize the service tool to collect data about specific IBM Spectrum Control components. For more information, see [How to customize the service tool](#).

Procedure

To run the service tool for all components, follow these steps:

1. Log on to the system where IBM Spectrum Control is installed.
2. Go to the following directory:

Windows operating systems:
installation_dir\service

Linux or AIX operating systems:
installation_dir/service/

3. Run the following program:

Windows operating systems:
service.bat

Linux or AIX operating systems:
service.sh

A compressed file, *SCServiceFiles_all.zip*, is created in the following directory:

Windows operating systems:
 `installation_dir\service\data\`
Linux or AIX operating systems:
 `installation_dir/service/data/`

4. Upload the log package by using one of the following methods:

If . . .	Then . . .
You are a US healthcare client	Blue Diamond data upload instructions.
The log file is < 200 MiB	Upload and attach the file to the case.
The log file is > 200 MiB	Enhanced Customer Data Repository (ECuRep) - Send data (FTP).

Related tasks

- [Creating a compressed file for a Storage Resource agent](#)

Related reference

- [How to customize the service tool](#)

Creating a compressed file for a Storage Resource agent

Run the service tool on Storage Resource agents that were deployed by using the web-based GUI to create a compressed file that can be sent to IBM® Software Support.

Before you begin

You must have administrator authority on Windows operating systems or root authority on AIX® and Linux® operating systems.

Procedure

To run the service tool on the Storage Resource agents, complete the following steps:

1. In the menu bar, go to Servers, > Servers.
2. Right-click the server where the Storage Resource agent is deployed, and select Logs, > Collect Agent Logs.

Results

The following compressed files are created:

Windows operating systems:
 `C:\Program Files\IBM\TPC\data\log\SRATraces\SRA_computer_name\SCServiceInfo.zip`
AIX and Linux operating systems:
 `/opt/IBM/TPC/data/log/SRATraces/SRA_computer_name/SCServiceInfo.zip`

where `SRA_computer_name` represents the name of the computer on which the Storage Resource agent is located.
If the compressed file cannot be created, a message indicates that the job was unsuccessful.

For more information about the error, see the server log file or the services script information file. The files are in one of the following default directories:

Server log file

This file is on the system where IBM Spectrum® Control is installed:

Windows operating systems:
 `c:\Program Files\IBM\TPC\data\log`
AIX or Linux operating systems:
 `/opt/IBM/TPC/data/log`

Services script information file

This file is on the computer on which the Storage Resource agent is installed:

Windows operating systems:
 `c:\Program Files\IBM\TPC\SRA_computer_name\services\SCServiceInfo.log`
AIX or Linux operating systems:
 `/opt/IBM/TPC/SRA_computer_name/services/SCServiceInfo.log`

For more information about customizing the data that is collected by the service tool, see [How to customize the service tool](#).

Related tasks

- [Packaging log files from the command line and sending them to IBM Support](#)

How to customize the service tool

The default behavior of the service tool is to collect data about all IBM Spectrum® Control components, but you can use the service tool to collect data about specific components. You can also use command-line parameters to specify a location to place the data that is collected, specify that the data is compressed, or to specify both.

Specifying help and output command-line parameters

To obtain information about the service tool usage, use the `-help` command-line parameter.

To specify the data that is collected by the service tool, use the following command-line parameters when you run service tool:

-output *directory_path*

Places the files that contain the data that was collected in a directory that you specify. If you specify a directory that does not exist on your system, that directory is created. If you do not use the `-output directory_path` parameter, the files are placed in the default directory:

Windows operating systems

`installation_dir\service\data`

Linux® or AIX® operating systems

`installation_dir/service/data`

Restriction: If you specify a directory, the directory path cannot contain spaces. This restriction refers to the `-output` option.

-pmr

The number of the PMR to which the support information is related. Use the following format to enter the number: `nnnnnn,nnn,nnn` where `n` represents a number, such as `12345,123,123`. If you enter a value for the `-pmr` parameter, the service tool automatically uses FTP to upload the service information to IBM®. If the service tool cannot upload the service information to IBM from the IBM Spectrum Control server, the generated file must be uploaded manually. The `-pmr` parameter is ignored if the `-nozip` parameter is used.

-nozip

When you use this parameter, the compressed collected data archives are no longer created. The service tool creates separate directories for each component for which data was collected. You can then create compressed archives for the collected files. In this way, you can control the size and content of each compressed file. To specify a directory other than the default directory, use the `-output directory_path` parameter.

Tip: You can specify more than one parameter, for example, `C:\Program Files\IBM\TPC\service>service -install -nozip`.

Specifying Javacore, Java heap dump, snap trace, and Java core dump file collection

By default, the service tool collects all Java™ core files, the most recent Java heap dump file, and the most recent snap trace file for the Data server, Device server, Alert server, web server, and data collector. The service tool can collect Java core dump files for these IBM Spectrum Control components, but it *does not* by default. You can edit the following parameters in the `installation_dir/service/service.properties` file in order to control which types of files and how many of each file type are collected by the service tool:

coreFiles=0

Specifies how many Java core dump files are collected. The default value is 0 because Java core dump files are generally very large. The valid values for this parameter are 0 or a positive integer. Any other value results in the default behavior.

javaCoreFiles=-1

Specifies how many Javacore txt files are collected. The default value of -1 means that all Javacore txt files are collected. The valid values for this parameter are 0 or a positive integer. Any other value results in the default behavior.

heapDumpFiles=1

Specifies how many Java heap dump files are collected. The default is to collect the most recent Java heap dump file. The valid values for this parameter are 0 or a positive integer. Any other value results in the default behavior.

snapFiles=1

Specifies how many snap trace files are collected. The default is to collect the most recent snap trace file. The valid values for this parameter are 0 or a positive integer. Any other value results in the default behavior.

Collecting data for specific IBM Spectrum Control components

You can use the service tool to collect data about specific IBM Spectrum Control components.

Use the following parameters to specify the components:

-all

All components. The default behavior is to collect data about all components.

-install

Installation component files.

-data

Data server component files.

-device

Device server component files.

-datacollector

Data collector component files.

-alert

Alert server component files.

-export

Export server component files.

-sra

Storage Resource agent component files.

-db

Db2® files.

-cli

Command-line interface files.

-gui

Collects data about GUI files.

If you collect information about a particular component, and you do not specify the `-nozip` parameter, you can identify the contents of a compressed service file from its name. For example, if you specify the `-db -gui` parameters but did not specify the `-nozip` parameter, a file named `SCServiceFiles_db_gui.zip` is created.

Warning: An existing compressed file is overwritten when another file of the same name is created. For example, if you run `C:\Program Files\IBM\TPC\service>service -db -gui`, a file named `SCServiceFiles_db_gui.zip` is created. If you rerun the tool with the same component options, a new file named `SCServiceFiles_db_gui.zip` is created. This new file overwrites the previously created file unless you specify the `-nozip` parameter, or use the `-output` parameter to specify a different path. If you specify the `-nozip` parameter, a directory is created for the components that you specified in the parameter. If you did not specify a specific component, data is collected for all installed components, and the data is placed in files in the following directory:

Windows operating systems:
`installation_dir\service\data`
Linux or AIX operating systems:
`installation_dir/service/data`

Files for particular components are then placed in a directory that corresponds to that component. Common files, such as `license.txt`, are placed in the following directory:

Windows operating systems:
`installation_dir\service\data`
Linux or AIX operating systems:
`installation_dir/service/data`

Restriction: You cannot specify a specific component as a command-line parameter when you specify the `-all` parameter. Also, when you specify a component that is not installed on the computer, the service tool displays an error message.

Related tasks

- [Troubleshooting FTP transfers](#)

Administering the IBM Spectrum Control database

The IBM Spectrum® Control database is the repository for information that is collected about the monitored resources in your environment.

- [Backing up the database](#)
Choose and then implement the Db2 backup method for securing the data that is collected and stored in the database for IBM Spectrum Control.
- [Restoring the database](#)
Choose and then implement the Db2 restore method for restoring your backup of the IBM Spectrum Control database.
- [Disaster recovery](#)
Back up your database regularly to be prepared for disaster recovery if a disaster event occurs.
- [Maintaining and improving the performance of the database](#)
You can use the database maintenance tool to calculate statistics for the IBM Spectrum Control database. The tool can also reorganize the database to restore efficiency and improve performance.
- [Repository copy tool](#)
You can use the Repository copy tool, **repocopy**, to export all the tables in the IBM Spectrum Control database repository for purposes of debugging problems.

Backing up the database

Choose and then implement the Db2® backup method for securing the data that is collected and stored in the database for IBM Spectrum® Control.

Backup types

To back up your database, choose one of the following options:

Offline (Default)
When the data is being backed up, you can neither access nor connect to the database.

Online
When the data is being backed up, you can access and connect to the database. Unlike the offline option, the database remains available to you and the applications that use the database. To configure this option, requires a good knowledge of Db2.

Output locations

To specify the location of the data that is backed up, choose one of the following options:

File system (Default)
Back up the data on a file system.
You can copy the file system that you specified to a removable tape or use IBM Spectrum Protect to back up the file system.

IBM Spectrum Protect
Back up the data to IBM Spectrum Protect. To back up the data, use IBM Spectrum Protect Backup/Archive client and client API on the same computer that hosts the IBM Spectrum Control Db2 databases.

Logging types

In the event of a system failure, the log files are used to recover data. You can choose either circular logging or archive logging.

The types of logging are:

Circular (Default)

This type of logging is used with IBM Spectrum Control for an offline backup.

Archive

This type of logging is used with an online backup. With archive logging, you can enable a rollforward recovery of the database to a specific date and time. A good knowledge of Db2 is needed to manage this type of logging.

- [Comparison of database backup methods](#)

The method that you choose to back up your data determines whether IBM Spectrum Control remains online or offline during the backup process.

- [Backing up the database offline using the command line](#)

By default, the IBM Spectrum Control database (TPCDB) is configured to use circular logging that requires backups to be performed offline. You can use the command line to perform the offline backup.

- [Backing up the database offline using IBM® Data Studio client](#)

By default, the IBM Spectrum Control database (TPCDB) is configured to use circular logging that requires backups to be performed offline. You can use the IBM® Data Studio client to perform the offline backup.

- [Backing up the database online using IBM® Data Studio client](#)

You can use the online method to back up a IBM Spectrum Control database and ensure continuous availability of the database and the applications that use it.

Related concepts

- [Comparison of database backup methods](#)

Related information

- [📖 Restore overview](#)
- [📖 Recover overview](#)
- [📖 IBM Redbook: IBM Tivoli Storage Productivity Center Beyond the Basics](#)

Comparison of database backup methods

The method that you choose to back up your data determines whether IBM Spectrum® Control remains online or offline during the backup process.

Advantages of an offline backup

The advantages of the offline backup method are:

- The offline backup method is the default method and it is easier than the online method to configure and to maintain.
- The circular type of logging that is used for offline backups is easier to configure and maintain than the type of logging that is used for online backups.

Disadvantages of an offline backup

The disadvantages of the offline backup method are:

- You must stop IBM Spectrum Control when you back up the data. So data is not collected and your storage resources are not being monitored during the back up process.
- You cannot collect performance data for the disk subsystems and SAN fabrics when data is being backed up.
- You might miss critical events, for example, failures within a SAN fabric, that occur during the backup process.

Tip: To minimize the loss of data for your storage resources and to ensure that you do not miss critical events, back up your data when your storage resources are not being used or when storage usage is low.

Advantages of an online backup

The advantages of the online backup method are:

- You continue to collect data and monitor your storage resources during the backup process because you do not have to stop IBM Spectrum Control.
- You continue to receive alerts and can respond quickly to critical events at any time of day.
- You continue to collect performance data for your disk subsystems and SAN fabrics.

Disadvantages of an online backup

The disadvantages of the online backup method are:

- The archive type of logging that is used with this type of backup is a more advanced method; it requires a good knowledge DB2® operation and administration.
- Software upgrades to IBM Spectrum Control that involve changes to the layout of the database might not complete successfully. In such cases, you can use circular logging to ensure that the software upgrade succeeds. You can switch back to archive logging after the software upgrade is installed.

Related information

- [📖 IBM DB2 11.5 for Linux, Unix and Windows](#)

- [IBM Redbook: IBM Tivoli Storage Productivity Center Beyond the Basics](#)

Backing up the database offline using the command line

By default, the IBM Spectrum Control database (TPCDB) is configured to use circular logging that requires backups to be performed offline. You can use the command line to perform the offline backup.

Procedure

To back up the database using the command line, complete the following steps:

1. Close the IBM Spectrum Control GUI.
2. Stop the IBM Spectrum Control services.
3. Complete these steps to initialize the Db2® environment:
 - a. While you are logged in to the Windows operating system as an Administrator, from the Windows Start menu, select IBM DB2 DB2COPY1 (Default) > DB2 Command Window - Administrator.
 - b. While you are logged in to the Linux® or AIX® operating system as the root user, switch to the user that is the Db2 instance owner (for example, db2inst1)
4. In the Db2 environment, run the following commands to prevent all users and applications from accessing Db2:

```
db2 force application all
db2 terminate
db2 list applications
```

5. Create a directory to store the backup of the IBM Spectrum Control database.
Tip: Choose a directory location that has enough free space to hold the number of backups that you plan to retain. Use a separate file system rather than the file system that contains the IBM Spectrum Control database. You can choose to use a location that is a remotely mounted Common Internet File System (CIFS) or Network File System (NFS), so the backup data is secured to another server.
6. In the Db2 environment, run the following command to backup the IBM Spectrum Control database:

```
DB2 BACKUP DATABASE TPCDB USER user_name USING password TO location COMPRESS
```

where *user_name* is the user who owns the Db2 instance where the IBM Spectrum Control database is located, *password* is the password that is associated with that user name, and *location* is the directory (created in step 5) where the backup is stored.

Examples:

- Windows operating system:

```
DB2 BACKUP DATABASE
TPCDB USER johndoe USING password1234 TO C:\DB_Backup COMPRESS
```
- Linux and AIX operating systems:

```
DB2 BACKUP DATABASE TPCDB USER johndoe USING password1234 TO
/tmp/DB_Backup COMPRESS
```

7. Restart the IBM Spectrum Control services.

Results

The offline database backup to a file system is run and the IBM Spectrum Control services are started again.

Related reference

- [BACKUP DATABASE command](#)
- [Stopping the IBM Spectrum Control servers by using scripts](#)
- [Starting the IBM Spectrum Control servers by using scripts](#)

Related information

- [IBM Data Studio documentation](#)
- [IBM DB2 11.5 for Linux, Unix and Windows](#)
- [IBM Redbook: IBM Tivoli Storage Productivity Center Beyond the Basics \(Chapter 3\)](#)

Backing up the database offline using IBM® Data Studio client

By default, the IBM Spectrum Control database (TPCDB) is configured to use circular logging that requires backups to be performed offline. You can use the IBM® Data Studio client to perform the offline backup.

Procedure

1. Stop the IBM Spectrum Control services.
2. Start the IBM Data Studio client.
To start the IBM Data Studio client from a command window, enter:

- On Windows operating system:

```
product_installation_directory\eclipse.exe
```

- On Linux® operating system:

`.product_installation_directory/eclipse`

3. On the New Connection page, select DB2 for Linux, UNIX and Windows and configure your IBM Spectrum Control database connection parameters.
4. Click Next.
5. On the Database Administration - IBM Data Studio page, expand DB2 .
6. Right-click DB2 and select TPCDB, Back Up and Restore, Back Up.
7. On the Back up TPCDB page, confirm the details of your database.
8. Click Backup Type.
9. Click Backup Image and select File System as the media type.
10. Click Backup Options.
11. In the Backup options compression and throttle section, select the Compress backup image and Throttle this utility to regulate the performance impact on the database workload check boxes.
12. Click Backup Performance.
To accelerate the backup process, you can increase the number of table spaces and buffers.
13. After you have set all of the options, click Run or Preview Command to see the actual Db2® command that is run to backup the data.
14. Restart the IBM Spectrum Control services.

Results

The offline database backup to a file system is run and the IBM Spectrum Control services are started again.

Backing up the database online using IBM® Data Studio client

You can use the online method to back up a IBM Spectrum Control database and ensure continuous availability of the database and the applications that use it.

About this task

When you use the online backup method, Db2® does not clean up old archive log files. You need to have processes in place to clean up the old log files after a specific amount of time to prevent the system from filling up. You also need to plan for the amount of space you might need. The log space that is required for a IBM Spectrum Control database can grow larger than the Db2 database over period of time.

Procedure

To backup the database, complete these steps:

1. Open IBM® Data Studio client.
2. Start the IBM Data Studio client.

To start the product from a command window, enter:

- On Windows operating system:

`product_installation_directory\eclipse.exe`

- On Linux® operating system:

`.product_installation_directory/eclipse`

3. On the New Connection page, select Db2 for Linux, UNIX and Windows and configure your IBM Spectrum Control database connection parameters.
4. Click Next.
5. On the Database Administration - IBM Data Studio page, expand Db2 .
6. Right-click Db2 for Linux, UNIX and Windows and select TPCDB, Set Up and Configure, Configure Database Logging.
7. On the Configure Database Logging TPCDB page, click Logging Type and select Archive.
8. Click Logging Size and enter your log file size information. For example, 2500.
9. Click Log Location and enter the Db2 log path information.
10. On the Database Administration - IBM Data Studio page, expand Db2 .
11. Right-click Db2 and select TPCDB and select Back Up and Restore, Back Up.
12. Click Backup Image and select File System as the media type.
13. Click Backup Options.
14. In the Backup options compression and throttle section, select the Compress backup image and Throttle this utility to regulate the performance impact on the database workload check boxes.
15. Click Backup Performance.
You can increase the number of table spaces and buffers, from the default values, to improve performance.
16. After you set all of the options, click Run or Preview Command to see the actual Db2 command that is run as part of the backup.

Results

The online database backup to a file system is run and the backup is complete.

Tip: You can also perform an online backup of the IBM Spectrum Control databases to an IBM Spectrum Protect server. The significant difference between online and offline backup is the need to enable archive logging on the databases. When you use the online method, it provides many backup and recovery benefits at the expense of increased complexity in the database operation. Set up and test your Db2 to IBM Spectrum Protect integration before you implement the online method to the IBM Spectrum Protect output destination to verify the communication is working properly.

Related information

- [IBM Data Studio documentation](#)
- [IBM Db2 11.5 for Linux, Unix and Windows](#)

Restoring the database

Choose and then implement the Db2® restore method for restoring your backup of the IBM Spectrum® Control database.

- [Restoring the database using the command line](#)
You can use the command line to restore your backup of the IBM Spectrum Control database.
- [Restoring the database using IBM Data Studio client](#)
To restore the IBM Spectrum Control database (TPCDB), use IBM® Data Studio client.

Related information

- [IBM Data Studio documentation](#)
- [IBM DB2 11.5 for Linux, Unix and Windows](#)
- [Starting the IBM Data Studio full client](#)
- [SQL0294N SQLB_CONTAINER IN USE error when restoring database](#)

Restoring the database using the command line

You can use the command line to restore your backup of the IBM Spectrum® Control database.

About this task

Restriction: Do not restore the IBM Spectrum Control database backup from one version of IBM Spectrum Control into another version. For example, do not restore a backup from IBM Spectrum Control 5.3.4 into 5.3.7, or any other.

Procedure

1. Stop the IBM Spectrum Control services.
2. Complete these steps to initialize the Db2® environment:
 - a. While you are logged in to the Windows operating system as an Administrator, from the Windows Start menu, select IBM DB2 DB2COPY1 (Default) > DB2 Command Window - Administrator.
 - b. While you are logged in to the Linux® or AIX® operating system as the root user, switch to the user that is the Db2 instance owner (for example, db2inst1)
3. In the Db2 environment, run the following commands to prevent all users and applications from accessing Db2:

```
db2 force application all
db2 terminate
db2 list applications
```

4. In the Db2 environment, run the following command to restore your backup of the IBM Spectrum Control database:

```
DB2 RESTORE DATABASE TPCDB FROM location INTO TPCDB REPLACE EXISTING
```

where *location* is the directory where you stored the backup.

Examples:

- Windows operating system: **DB2 RESTORE DATABASE TPCDB FROM C:\DB_Backup INTO TPCDB REPLACE EXISTING**
- Linux and AIX operating systems: **DB2 RESTORE DATABASE TPCDB FROM /tmp/DB_Backup INTO TPCDB REPLACE EXISTING**

5. Restart the IBM Spectrum Control services.

Restoring the database using IBM Data Studio client

To restore the IBM Spectrum® Control database (TPCDB), use IBM® Data Studio client.

Before you begin

Download and install the IBM Data Studio client. For more information about how to download and install the IBM Data Studio client, see <http://www.ibm.com/developerworks/downloads/im/data/>.

Before you can use the IBM Data Studio client to restore the IBM Spectrum Control database, you must add a connection to the database in the IBM Data Studio client.

Restriction:

Do not restore the IBM Spectrum Control database backup from one version of IBM Spectrum Control into another version. For example, do not restore a backup from IBM Spectrum Control Version 5.3.4 into Version 5.3.7, or any other version.

Procedure

1. Stop the IBM Spectrum Control services.
2. Start the IBM Data Studio client.
3. Click the Administer tab in the Data Studio client Task Launcher, then click Connect and browse a database.

4. On the Connection Parameters page, click DB2 for Linux, UNIX, and Windows and configure your IBM Spectrum Control database connection parameters.
5. Click Next.
6. On the Database Administration - IBM Data Studio page, expand DB2.
7. Right-click TPCDB, then click Back Up and Restore, Restore.
8. On the Restore Database TPCDB page, click Restore backup to the current database to set the restore type.
9. Click Restore Objects, and then click Restore the entire database.
10. Select the backup image that you want to restore.
11. On the Restore Database TPCDB page, click Run.
12. Restart the IBM Spectrum Control services.

Disaster recovery

Back up your database regularly to be prepared for disaster recovery if a disaster event occurs.

Disaster recovery is the rebuilding of a database or table space after a disaster event such as media or storage failure, power interruption, or application failure occurs. If a disaster event occurs and a database or table space is damaged or corrupted, you can restore one of your backups.

Related tasks

- [Restoring the database](#)

Related reference

- [Backing up the database](#)

Related information

-  [Data recovery](#)

Maintaining and improving the performance of the database

You can use the database maintenance tool to calculate statistics for the IBM Spectrum Control database. The tool can also reorganize the database to restore efficiency and improve performance.

About this task

By default, the database maintenance tool runs the Db2® **runstats** command on all database tables that are used by IBM Spectrum Control. By using the **runstats** command, the database maintenance tool updates statistics about the characteristics of a table and its associated indexes. Because Db2 uses the statistics to determine access paths to data, when you run the database maintenance tool you help to ensure the effectiveness of the paths that are selected.

By specifying an option of the database maintenance tool, you can instruct the tool to reorganize the database tables, if necessary. A set of formulas are applied to the statistics that were collected about the database to determine if reorganization is necessary. Tables are reorganized by reconstructing rows to eliminate fragmented data and by compacting information. Index data is reorganized into unfragmented, physically contiguous pages.

You can customize the reorganization function in the database maintenance tool by updating a properties file. Properties determine which formulas can trigger database reorganization. You can specify properties that exclude tables from being reorganized based on size. Also, you can specify a property to force database reorganization.

- [Updating database statistics](#)
Use the database maintenance tool to update statistics about the databases that are used by IBM Spectrum Control. Because DB2® uses these statistics to select access paths to data, when you run the database maintenance tool, you can help to improve the effectiveness of the paths that DB2 selects.
- [Reorganizing database tables](#)
Use the database maintenance tool to analyze the database tables that are used by IBM Spectrum Control. If necessary, the database maintenance tool reorganizes the database tables and indexes.
- [Customizing the reorganization function of the database maintenance tool](#)
You can customize the database maintenance tool to specify which formulas determine whether a database is reorganized.

Updating database statistics

Use the database maintenance tool to update statistics about the databases that are used by IBM Spectrum® Control. Because DB2® uses these statistics to select access paths to data, when you run the database maintenance tool, you can help to improve the effectiveness of the paths that DB2 selects.

About this task

Run the database maintenance tool to make substantial changes to the IBM Spectrum Control database, such as numerous table space updates, deletions, or insertions. The database maintenance tool uses the **RUNSTATS** command to update statistics about the physical characteristics of a table and the associated indexes.

Note:

The database maintenance script for updating statistics for IBM Spectrum Control can be run online or offline.

Procedure

To update statistics for databases, follow these steps:

1. Log on to the computer where you installed IBM Spectrum Control.
2. Change the directory.

```
Linux® and UNIX operating systems
/opt/IBM/TPC/data/server/tools/
Windows operating systems
C:\Program files\IBM\TPC\data\server\tools\
```

3. Enter the following command to run the database maintenance tool:

```
Linux and UNIX operating systems
runTPCDBMaintenance
Windows operating systems
runTPCDBMaintenance.bat
```

Note: To monitor the progress of a statistical update for the database, redirect your output to a file.

Example

```
runTPCDBMaintenance > /tmp/tpcdb_update_stats.txt
```

Reorganizing database tables

Use the database maintenance tool to analyze the database tables that are used by IBM Spectrum® Control. If necessary, the database maintenance tool reorganizes the database tables and indexes.

About this task

The database maintenance tool uses a set of formulas to analyze the physical location of rows, and the size of the tables; it analyzes the indexes and their relationship to the table. If the calculated result of a formula exceeds set boundaries, the tool reorganizes the tables and indexes as needed. The tool reorganizes database tables, if necessary, by reconstructing rows to eliminate fragmented data. The tool reorganizes index data, if necessary, into unfragmented, physically contiguous pages.

Tips:

- Run the database maintenance script for analyzing database tables while IBM Spectrum Control is offline; no database workload is present.
- Ensure that you have enough available capacity to reorganize the database tables. An offline reorg might require an amount of available capacity that is to 2 to 3 times the size of the existing IBM Spectrum Control tables.

Procedure

To analyze and reorganize databases, follow these steps:

1. Log on to the computer where you installed IBM Spectrum Control.
2. Change the directory.

```
Linux® and UNIX operating systems
/opt/IBM/TPC/data/server/tools/
Windows operating systems
C:\Program files\IBM\TPC\data\server\tools\
```

3. Enter the following command to run the database maintenance tool to analyze and reorganize databases:

```
Linux and UNIX operating systems
runTPCDBMaintenance reorg
Windows operating systems
runTPCDBMaintenance.bat reorg
```

Example

```
runTPCDBMaintenance reorg > /tmp/tpcdb_reorg.txt
```

Note: To monitor the progress of the database reorganization, redirect your output to a file.

Customizing the reorganization function of the database maintenance tool

You can customize the database maintenance tool to specify which formulas determine whether a database is reorganized.

About this task

To specify which formulas can trigger a database reorganization, edit the properties in the TPCDBMaintenance.properties file. For example, you can customize the tool to ignore tables that are smaller or larger than the defined size limits, or customize it to always reorganize the database.

Procedure

To analyze and reorganize databases, follow these steps:

1. Log on to the system where IBM Spectrum® Control is installed.
2. Change to the following directory:

```
Linux® and UNIX operating systems
/opt/IBM/TPC/data/server/tools/
Windows operating systems
C:\Program files\IBM\TPC\data\server\tools\
```

3. Open TPCDBMaintenance.properties in a text editor and modify the property settings as needed.

The f1-f8 properties all refer to the same formulas that are used by the DB2® **REORGCHK** command. For more information about any of these formulas, see the DB2 product documentation about the **REORGCHK** command at <https://www.ibm.com/docs/en/db2>.

You can modify the following properties:

f1= {true | false}
Specifies whether the result of formula 1 can trigger a database table reorganization. Formula 1 checks the number of overflow rows in a table.

f2= {true | false}
Specifies whether the result of formula 2 can trigger a database table reorganization. Formula 2 checks the effective space utilization of data pages.

f3= {true | false}
Specifies whether the result of formula 3 can trigger a database table reorganization. Formula 3 checks the number of empty pages. Pages can become empty after rows are deleted.

f4= {true | false}
Specifies whether the result of formula 4 can trigger the reorganization of index data. Formula 4 checks the clustering ratio of an index.

f5= {true | false}
Specifies whether the result of formula 5 can trigger the reorganization of index data. Formula 5 checks the space that is reserved for index entries.

f6= {true | false}
Specifies whether the result of formula 6 can trigger the reorganization of index data. Formula 6 determines whether re-creating an index would result in a tree with fewer levels.

f7= {true | false}
Specifies whether the result of formula 7 can trigger the reorganization of index data. Formula 7 checks the number of pseudo-deleted RIDs on non-pseudo-empty pages.

f8= {true | false}
Specifies whether the result of formula 8 can trigger the reorganization of index data. Formula 8 checks the number of pseudo-empty leaf pages.

maxReorgTableSize= {size_in_bytes | none}
Specifies the maximum size that a database table must be to be considered for reorganization.

minReorgTableSize= {size_in_bytes | none}
Specifies the minimum size that a database table must be to be considered for reorganization.

forceReorg= {true | false}
Specifies whether the database is always reorganized by the database maintenance tool when the reorg argument is specified. If this property is set to true, all other properties in the file are ignored.

4. Save TPCDBMaintenance.properties.

Repository copy tool

You can use the Repository copy tool, **repocopy**, to export all the tables in the IBM Spectrum® Control database repository for purposes of debugging problems.

You can send the exported data to IBM® Software Support to help debug problems.

Tip: If you want to export only performance data from the IBM Spectrum Control repository, you can create performance support packages. You can create performance support packages for storage systems or fabrics. For more information about exporting performance support packages, see [Exporting performance data for storage systems and fabrics](#).

- [Exporting repository data](#)
Use the Repository copy tool to export data from an existing repository into a text file.

Exporting repository data

Use the Repository copy tool to export data from an existing repository into a text file.

Procedure

To export repository data, follow these steps:

1. Go to the following default directory:

```
Windows operating systems:
c:\Program Files\IBM\TPC\data\server\tools
```

```
Linux® or AIX® operating systems:
/opt/IBM/TPC/data/server/tools
```

2. Issue the **repocopy** command:

```
Windows operating systems:
repocopy.bat
```


Linux or AIX operating systems:

repocopy

3. Select Export data from repository tables and click Next.
4. In the Options for Import/Export window, enter information in the following fields:

Directory for Export

Enter the directory where the comma-delimited file is saved.

Delimiter

Enter a delimiter for the delimited file format (a comma is the default).

Quote

Enter the symbol that contains string data (double quotation marks is the default).

IBM Spectrum® Control exports the data into the comma-delimited file that you specify, and places it in a file named *tablename.txt*. Click Next.

5. Select one of the following options and click Next.
 - Export by using DB2® native format.
 - Export by using text files (the preferred method).
6. Select one of the following options and click Next.
 - Export base tables (always export the base tables)
 - Export Performance Manager tables, if requested by IBM® Software Support
 - Export history tables, if requested by IBM Software Support.

The information that is detected in the server.config file is displayed in the Connection Properties window within the following fields:

- Database Types
- User name
- Password
- Driver Class
- Driver URL
- Database
- DB Creator
- Classpath

If you want to export data from a different database from the one listed in the server.config file, you can select the database from the Database Types list box.

Manually enter the database information.

7. Click Finish.
8. Click Run.

Results

As you progress through the export process, messages are written to a progress log that is displayed. You can track the steps through the progress log.

When the **repocopy** command is used with a remote database, the DB2 shared library is not available for loading the **libTSRMinsubd.so** file. You can ignore this message. Click OK and continue.

Administering Db2

Administer IBM® Db2® by backing up the IBM Spectrum® Control database, starting the IBM Data Studio full client, and starting and stopping Db2.

- [Using the command line on UNIX and Linux](#)
This topic describes how to use a command line to perform actions against a IBM Db2 instance under UNIX or Linux®.
- [Manually starting Db2 on Windows](#)
Start DB2® on Windows operating systems.
- [Manually stopping Db2 on Windows](#)
You can stop IBM Db2 manually on Windows.
- [Starting the IBM Data Studio full client](#)
You can start the Data Studio full client on your workstation from either a menu option or the command line.
- [Monitoring Db2](#)
The minimum user authority level needed for monitoring IBM Db2 instances is a user with Db2® system maintenance authority (SYSMAINT).

Related reference

- [Db2 and database troubleshooting](#)

Using the command line on UNIX and Linux

This topic describes how to use a command line to perform actions against a IBM® Db2® instance under UNIX or Linux®.

About this task

If the Db2 Control Center is unavailable or you do not have access to a graphical user interface, you can use a command line to execute Db2 commands such as starting and stopping an instance.

Important: If you are using DB2® 10.1 or higher, you must use the command-line interface.

To use a command line to perform actions against an instance of Db2, complete the following steps:

Procedure

1. Log in with a user ID or name that has ROOT, SYSADM, SYSCTRL, or SYSMAINT authority on the instance; or log in as the instance owner.
2. Run the startup script:
 - For Bourne or Korn shell, type: **. HOME/sqllib/db2profile**.
 - For C shell, type: **source HOME/sqllib/db2cshrc**.where **HOME** is the home directory of the instance you want to use.
3. To start the instance using the command line, type **db2start**.
Note: When you run commands to start the database manager instance, the Db2 database manager applies the command to the current instance.
4. To stop the instance using the command line, type **db2stop**.
Note: When you run commands to stop the database manager instance, the Db2 database manager applies the command to the current instance.

Manually starting Db2 on Windows

Start DB2® on Windows operating systems.

About this task

To start Db2® manually, complete the following steps:

Procedure

1. Start the following Windows services:
 - DB - DB2-0
 - DB2DAS - DB2DAS00
 - DB2 JDBC Applet Server
 - DB2 License Server
 - DB2 Security Server
2. Open a Db2 command window.
3. From the Db2 Command window, run the **db2start** command.

Related information

- [Accessing administration tools](#)
- [Opening IBM Spectrum Control GUIs and CLIs](#)

Manually stopping Db2 on Windows

You can stop IBM® Db2® manually on Windows.

Procedure

1. Stop the following Windows services:
 - Db2 Security Server
 - Db2 License Server
 - Db2 JDBC Applet Server
 - Db2DAS - Db2DAS00
 - Db2 - Db2-0Note: When you stop the Db2 Security Server service, you are prompted to stop the Warehouse logger and Warehouse Serve. Click Yes.
2. To open a Db2 command window, click Start, > Programs, > IBM Db2, > Command Line Tools, > Command Window.
3. From the Db2 Command window, issue the **db2stop** command.

Results

To restart Db2, issue the **db2start** command from the Db2 Command window. Before you can issue the command, you must first start these services on the Windows Services panel:

```
DB2 - DB2-0
DB2DAS - DB2DAS00
DB2 JDBC Applet Server
DB2 License Server
DB2 Security Server
```

Related tasks

- [Accessing administration tools](#)
- [Opening IBM Spectrum Control GUIs and CLIs](#)

Starting the IBM Data Studio full client

You can start the Data Studio full client on your workstation from either a menu option or the command line.

Opening IBM Data Studio Administration client on Windows operating systems

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	a. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start. b. On the Start page, right-click, and then click the All apps taskbar button.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start > All Programs.

2. Click IBM Data Studio > Data Studio Administration Client.

Open IBM Data Studio Administration client on Linux® and AIX® operating systems

On the command line, issue the following command:

```
DS_install_dir/eclipse
```

where DS_install_dir is the directory where you installed the full client.

Example

```
/opt/IBM/DS3.1.1/eclipse
```

Monitoring Db2

The minimum user authority level needed for monitoring IBM® Db2® instances is a user with Db2® system maintenance authority (SYSMAINT).

To check and set SYSMAINT authority, follow these steps:

1. Run this command in the Db2 command prompt window to check to see if there is an operating system user group defined to have SYSMAINT authority:

```
db2 get dbm cfg
```

In the output file, look for this information:

```
SYSADM group name      (SYSADM_GROUP) =
SYSCtrl group name     (SYSCtrl_GROUP) =
SYSMAINT group name    (SYSMAINT_GROUP) =
SYSMON group name      (SYSMON_GROUP) =
```

If the setup for the operating system group has not been done, you do not see a value set.

If the setup has been done, this example shows what you can expect to see:

```
SYSADM group name      (SYSADM_GROUP) =
SYSCtrl group name     (SYSCtrl_GROUP) =
SYSMAINT group name    (SYSMAINT_GROUP) = ADMINISTRATORS
SYSMON group name      (SYSMON_GROUP) =
```

In this example, the "ADMINISTRATORS" group has SYSMAINT_GROUP authority.

2. If the setup has been done, add the user you want to use to the ADMINISTRATORS group using the operating system utilities or use a user that already belongs to the ADMINISTRATORS group.

If you want to give a user group "SYSMAINT_GROUP" authority, follow these steps:

- a. If a user (for example **userA**) belongs to an operating system group called **db2monitor**, here is an example of setting the **db2monitor** group with SYSMAINT authority. From the Db2 command prompt window, run the following command:

```
db2 update dbm cfg using SYSMAINT_GROUP db2monitor
```

- b. After issuing the **db2 update** command, restart Db2 by running the following command from the Db2 command prompt window or restarting the system:

```
db2 force application all
```

This command might need to be issued a few times to stop all the database connections.

- c. Run the following commands from the Db2 command prompt window:

```
db2stop
db2start
```

- d. **UserA** can now monitor the Db2 database.

Managing resources

Use IBM Spectrum® Control to monitor and manage the resources and storage infrastructure in an enterprise environment. In the GUI, you can view the condition, capacity, performance, and relationships of storage resources.

- **Resources that you can monitor**

You can use IBM Spectrum Control to monitor storage systems, servers, hypervisors, fabrics, and SAN switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.

- [Adding resources](#)
You must add resources to IBM Spectrum Control before you can collect data, generate reports, and manage storage that is related to those resources.
- [Adding rollup servers](#)
Rollup servers gather capacity and status information from IBM Spectrum Control servers in your enterprise. You can add rollup servers for monitoring on the Settings > Rollup Server Connections page.
- [Removing resources](#)
Remove resources that you no longer want to monitor with IBM Spectrum Control.
- [Collecting data](#)
Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. IBM Spectrum Control provides two different jobs for collecting data about resources.
- [Alerting](#)
Specify conditions that trigger alerts and the actions to take when those alerts are triggered, such as notify an email address. Use alert policies to define those alert conditions and notification settings for a group of resources.
- [Viewing information about resources](#)
You can view detailed information about the resources that are monitored by IBM Spectrum Control. Resources include storage systems, volumes, pools, servers, hypervisors, fabrics, and switches. You can also view information about internal and related resources.
- [Monitoring the status and condition of resources](#)
Monitor the operational condition of storage systems, servers, hypervisors, fabrics, and switches and the status of their internal resources. Use this information to identify potential problem areas in a storage environment.
- [Monitoring the performance of resources](#)
IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes key performance metrics that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.
- [Monitoring the capacity of resources](#)
IBM Spectrum Control can collect information about the capacity and space usage of block, file, and object storage resources. This information includes key metrics that can help you measure, identify, and troubleshoot capacity and space usage issues in your storage. You can view capacity metrics for storage resources, including tiers.
- [Monitoring and administering applications, departments, and general groups](#)
You create applications to monitor the storage capacity, space usage, and performance of applications, and you create departments to monitor the space usage of the applications in the department. You can structure your applications and departments hierarchically to match the structure of your business organization.
- [Provisioning storage](#)
You can assign storage to servers, hypervisors, and clusters on servers and hypervisors in the GUI.
- [Optimizing storage tiering](#)
To optimize the placement of volumes on storage tiers, analyze the tiering of volumes in your storage environment.
- [Optimizing storage pools](#)
Analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.
- [Reclaiming storage](#)
Use the recommendations to reclaim capacity before you plan new capacity purchases.
- [Transforming and migrating volumes](#)
You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier®.
- [Modifying the period for analyzing performance data](#)
To change the default period for analyzing performance data in storage virtualizer pools, issue the `setdscfg` command. You can extend or shorten the default period for a more accurate analysis.
- [Modifying the properties of resources](#)
Add or change the properties for individual resources or for multiple resources. You can use the properties to filter or sort the resources in the GUI or in an external application if the data is shared or exported.
- [Opening the management GUI for a resource](#)
Resources such as storage systems and switches can have their own management GUIs or element managers. In IBM Spectrum Control, you can open the start page in these GUIs from resource list pages, such as the Storage Systems page and Switches page.
- [Exporting information to a file](#)
You can save information about resources, tasks, or alerts to a PDF, CSV, or HTML file. Information that you can export from the GUI includes all the values that are being shown in the columns for a list of resources, tasks, or alerts.
- [Customizing lists](#)
Customize lists to focus on the information that is important to you. You can filter information, sort rows, and show, hide, and reorder columns in the table views for resources, tasks, alerts, and other objects.
- [Managing tasks](#)
Tasks are used to provision storage and optimize resources in your storage environment. Use the Tasks page to manage all the tasks that are used by IBM Spectrum Control to provision storage and optimize resources. Use resource list and resource details pages to manage the tasks for specific resources and resource types.
- [Planning copy data resources](#)
To plan for your block storage requirements, you need to be able to see how you currently use block storage in your storage environment.
- [IBM Spectrum Control REST API](#)
You can use the Representational State Transfer (REST) API for IBM Spectrum Control to access information about resources and to generate custom capacity, configuration, and performance reports.

Related tasks

- [Starting IBM Spectrum Control](#)

Related reference

- [Troubleshooting and problem determination](#)

Resources that you can monitor

You can use IBM Spectrum® Control to monitor storage systems, servers, hypervisors, fabrics, and SAN switches. Information about these top-level resources includes information about their internal resources and related resources. Internal resources are components that exist in a top-level resource. Related resources are external to a top-level resource, but are related to it through assigned storage, a network connection, or virtual hosting.








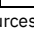








Information about object storage systems includes information about their object resources and related resources. Object resources are resources in an IBM Spectrum Scale cluster, such as accounts and containers, that enable the cluster to be used as an object storage system.

Icon legend: Each resource that you monitor with IBM Spectrum Control is represented by an icon for easy identification. View the following information to see the icons that are associated with each resource.

Fabrics

A *fabric* is a network of hubs, switches, adapter endpoints, and connecting cables that support a communication protocol between devices. The following table shows the internal resources and related resources that you can monitor for fabrics.

Table 1. Fabric resources that you can monitor


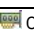





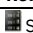



Fabric icon	Internal resources	Related resources	*Connected resources (to ports)
 Physical fabric  Virtual fabric	 Logical switches  Physical switches  Inter-switch connections  Blades  Ports  Zone sets	 Hypervisors  Servers  Storage systems  NPV switches (Cisco switches in NPV mode and Brocade switches in Access Gateway mode)  Discovered ports	All monitored resources that are connected to ports, plus the following resources:  Unmanaged resource. The name is determined from a fabric zone alias.  Unmanaged resource. The name is determined from a storage system's mapping to a host.  Unmanaged resource. The name is determined from one of a storage system's mappings to a host when multiple mappings are found.
*Connected resources (to ports): The icons for any resources that are managed by IBM Spectrum Control and connected to a port are shown in this column. Icons are also shown for connected resources whose names are known from a fabric zone alias or a storage system's mapping to a host (host connection) but are not currently being managed. If both a fabric zone alias and a storage system's mapping to a host are found for the same unmanaged resource, the fabric zone alias is shown as the name.			

Hypervisors

A *hypervisor* is a layer of software or a physical device that manages multiple, virtual instances of operating systems on the same hardware. Hypervisors that you can monitor include VMWare vSphere data sources such as ESX, ESXi, and vCenter servers. If a hypervisor is part of a cluster, information about that cluster is also available.

The following table shows the internal resources and related resources that you can monitor for hypervisors.


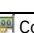
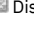

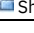

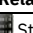


Table 2. Hypervisor resources that you can monitor

Hypervisor icon	Internal resources	Related resources
	 Controllers  Disks  Data stores  Virtual machine that is associated with an agentless server  Virtual machine that is associated with a Storage Resource agent  Virtual machine that was discovered, but isn't associated with an agentless server or Storage Resource agent	 Storage systems  Fabrics  Switches  Cluster

Servers

A *server* is a computer or host that is connected to a network and provides an access point to that network. If a server is part of a cluster, information about that cluster is also available. You can monitor the following internal resources and related resources for servers.

Table 3. Server resources that you can monitor

Server icon	Internal resources	Related resources
	 Controllers  Disks  Volume groups  File systems and logical volumes  Shares and exports	 Storage systems  Fabrics  Switches

Storage systems


A *storage system* is a device, such as a Redundant Array of Independent Disks (RAID) controller that creates and manages other storage devices. A storage virtualizer is a storage system that virtualizes storage space from internal storage or from another storage system.

Storage systems can be configured for block storage, file storage, a combination of block and file storage, or a combination of file and object storage. The method that a storage system uses for managing data determines the internal resources or object resources that are monitored.

The following tables show the internal resources, object resources, and related resources that you can monitor for block, file, and object storage systems.


Block storage systems

Table 4. Block storage system resources that you can monitor

Block storage system icon	Internal resources	Related resources
	<ul style="list-style-type: none"> Active quorum Disks Drives Enclosures External disks FC ports Host connections I/O groups IP ports Managed disks Modules Nodes Pools RAID arrays Volumes 	<ul style="list-style-type: none"> Back-end storage systems Fabrics Hypervisors Servers Switches Virtualizer storage systems GPFS clusters

File storage systems

Table 5. File storage system resources that you can monitor

File storage system icon	Internal resources	Related resources
	<ul style="list-style-type: none"> Disk Controllers Clusters Disks Shares Filesets File systems Network shared disks Nodes Pools Snapshots Quotas 	<ul style="list-style-type: none"> Block storage systems Fabrics GPFS clusters Object storage systems Hypervisors Servers Switches

Object storage systems












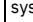




Table 6. Object storage system resources that you can monitor

Object storage system icons	Object internal resources	Object resources	Related resources
<ul style="list-style-type: none"> Object storage system IBM® Cloud Object Storage 	<ul style="list-style-type: none"> Access pools Accesser® nodes Mirrors Sites Slicestor® nodes Storage pools Vaults 	<ul style="list-style-type: none"> Accounts Containers 	<ul style="list-style-type: none"> Back-end storage systems Fabrics GPFS clusters Hypervisors Servers Switches

Switches

A SAN switch is a network infrastructure component to which multiple nodes attach. In IBM Spectrum Control, you can monitor Fibre Channel (FC) switches. The following table shows the internal resources and related resources that you can monitor for switches.

Table 7. Switch resources that you can monitor

Switch icon	Internal resources	Related resources	*Connected resources (to ports)
 Physical switch  Logical switch	 Inter-switch connections  Blades  Logical switches  Physical switches  Ports	 Fabrics  Hypervisors  Servers  Storage systems  Switches  Discovered ports	All monitored resources that are connected to ports, plus the following resources:  Unmanaged resource. The name is determined from a fabric zone alias.  Unmanaged resource. The name is determined from a storage system's mapping to a host.  Unmanaged resource. The name is determined from one of a storage system's mappings to a host when multiple mappings are found.

*Connected resources (to ports): The icons for any resources that are managed by IBM Spectrum Control and connected to a port are shown in this column. Icons are also shown for connected resources whose names are known from a fabric zone alias or a storage system's mapping to a host (host connection) but are not currently being managed. If both a fabric zone alias and a storage system's mapping to a host are found for the same unmanaged resource, the fabric zone alias is shown as the name.

Related tasks

- [Adding ESX and ESXi hypervisors](#)
- [Adding storage systems](#)
- [Adding fabrics and switches](#)

Adding resources

You must add resources to IBM Spectrum® Control before you can collect data, generate reports, and manage storage that is related to those resources.

Before you begin

Restriction: You can't collect performance metadata about a device that includes a forward slash (/) or backslash (\) in its name. Before you add a device for monitoring, remove any slashes in its name.

- [Required user roles for monitoring resources](#)
When you add storage systems, switches, and hypervisors for monitoring, you must provide a user name and password for logging in to those resources. The role or user group that is assigned to the user name determines the data collection and storage functions that you can use.
- [Adding storage systems](#)
Add the storage systems that you want to monitor in IBM Spectrum Control.
- [Adding fabrics and switches](#)
Add fabrics and switches for monitoring.
- [Adding hypervisors](#)
You can add hypervisors for monitoring by IBM Spectrum Control. Hypervisors include VMware vSphere data sources such as ESX and ESXi, and vCenter Server systems. When you add vSphere data sources, you can then collect data, generate reports, and manage storage that is related to those resources.
- [Adding servers](#)
Add servers for monitoring. IBM Spectrum Control creates and updates agentless servers automatically after it probes storage systems and hypervisors. You can also add a server by deploying a Storage Resource agent.

Related concepts

- [Alerting](#)

Related reference

- [Triggering conditions for alerts](#)

Required user roles for monitoring resources

When you add storage systems, switches, and hypervisors for monitoring, you must provide a user name and password for logging in to those resources. The role or user group that is assigned to the user name determines the data collection and storage functions that you can use.

The following roles are associated with the user names that IBM Spectrum® Control uses to log in to resources. Specify user names when you add a resource for monitoring. These roles are different from IBM Spectrum Control roles, which are assigned to users that log in to IBM Spectrum Control. For more information about IBM Spectrum Control roles, see [Role-based authorization](#).

Any roles that are not listed, but include the privileges of the roles that are listed, can also be used for monitoring resources.

Table 1. Required roles for storage system, switch, and hypervisor users

Resource	Required role for performance monitors	Required role for probes	Required role for provisioning, optimizing, and transforming
DS8000®	Monitor role or higher	Monitor role or higher	Administrator, Physical operator, or Logical operator
SAN Volume Controller	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For 8.3.1.2 or later, any role but some limitations might apply. Learn more	Monitor role or higher	Administrator
Storwize® V3500 Storwize V3700 Storwize V5000 Storwize V7000	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For 8.3.1.2 or later, any role but some limitations might apply. Learn more	Monitor role or higher	Administrator
Storwize V7000 Unified (block storage)	Administrator	Monitor role or higher	Administrator
FlashSystem 5000 FlashSystem 5100 FlashSystem 7200 FlashSystem 9100 FlashSystem 9200 FlashSystem V9000	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For 8.3.1.2 or later, any role but some limitations might apply. Learn more	Monitor role or higher	Administrator
IBM Spectrum Virtualize for Public Cloud	For versions earlier than 8.3.1.2, Administrator, or SecurityAdmin. For 8.3.1.2 or later, any role but some limitations might apply. Learn more	Monitor role or higher	Administrator
Storwize V7000 Unified (file storage)	Performance monitoring is not available	Monitor	Administrator, or Export Administrator + Storage Administrator.
XIV®	Monitor role or higher	Any role	Storage administrator
IBM Spectrum Accelerate IBM FlashSystem® A9000 IBM FlashSystem A9000R	Monitor role or higher	Any role	Not available
IBM FlashSystem 900	Monitor role or higher	Monitor role or higher	Not available
IBM Spectrum Scale (file storage)	Any role	Root or non-root*	Not available
IBM Spectrum Scale (object storage)	Performance monitoring is not available	The Keystone admin role. Information is collected only for the object storage accounts and containers that the user has access to. If you want to monitor all accounts and containers, the user must also be assigned the role that is defined in the reseller_admin_role configuration option in the Swift proxy server. The default value for the reseller_admin_role option is ResellerAdmin.	Not available
IBM® Cloud Object Storage	Performance monitoring is not available	Operator, System Administrator, or Super User role	Not available
Dell EMC storage systems	Operator role or higher	Operator role or higher	Not available

Resource	Required role for performance monitors	Required role for probes	Required role for provisioning, optimizing, and transforming
Hitachi storage systems	In Hitachi Device Manager: The user for the Export Tool must have the Storage Administrator (Performance Management) role	In Hitachi Command Suite: ViewGroup or higher and AdminGroup permission.	Provisioning is not available
NetApp storage systems	Operator role or higher	Operator role or higher	Not available
Pure storage systems	Role with read-only permission or higher	Role with read-only permission or higher	Not available
Cisco switches	Network-admin	Network-admin	Not applicable
Brocade switches with Fabric OS 8.2.1 or later	User or admin role that has the chassis-role permission	User or admin role that has the chassis-role permission	Not applicable
Brocade switches with a Fabric OS version earlier than 8.2.1	Administrator user in Brocade Network Advisor (BNA)	Administrator user in BNA	Not applicable
Hypervisors (such as ESX, ESXi, and vCenter Servers)	Performance monitoring is not available	A role that has permission to browse through data stores. For example: Administrator role or Virtual Machine Power® User role. For more information about roles and permission to browse data stores, see Checking permissions to browse data stores .	Any role
Rollup servers	Performance monitoring is not available	Administrator	Not available
Note: * You can add IBM Spectrum Scale and GSS systems as a non-root user, but that user must have privileges to run a set of specified administration commands using the sudo command on the cluster node. For more information, see Monitoring IBM Spectrum Scale without requiring root privileges .			

Adding storage systems

Add the storage systems that you want to monitor in IBM Spectrum® Control.

About this task

When you add block storage systems, asset, capacity, and configuration metadata and performance metadata is automatically collected and analyzed. When you add storage systems that manage block and file storage, the asset, capacity, and configuration metadata for file storage is also automatically collected and analyzed. Metadata Collection Schedule: By default, asset, capacity, and configuration metadata is aggregated and collected daily. Performance metadata is collected at the following intervals:

- Every 1 minute for IBM® block storage systems
- Every 5 minutes for Dell EMC Unity storage systems
- Every 5 minutes for Hitachi VSP storage systems
- Every 5 minutes for NetApp storage systems that are running ONTAP 9
- Every 5 minutes for Pure storage systems
- Every 15 minutes for other Dell EMC block storage systems.

You can add multiple storage systems at the same time, or add them one by one. To add multiple storage systems, the storage systems must be of the same type and share authentication credentials.

Important: To add storage systems, you must have Administrator privileges.

You can add the following types of storage systems for monitoring:

Table 1. Storage systems that can be monitored in IBM Spectrum Control

Storage system	Can add many at same time.
DS8000®	✓
Dell EMC Unity	✓
Dell EMC VMAX	
Dell EMC VNX	
Dell EMC VNXe	
FlashSystem 5000	✓
FlashSystem 5100	✓
FlashSystem 5200	✓
FlashSystem 7200	✓
FlashSystem 7300	✓
FlashSystem 9100	✓
FlashSystem 9200	✓
FlashSystem 9500	✓
FlashSystem V9000	✓

Storage system	Can add many at same time.
FlashSystem 900	✓
FlashSystem A9000	✓
FlashSystem A9000R	✓
Hitachi VSP	✓
NetApp ONTAP 9	✓
NetApp ONTAP 8 (7- mode)	
Pure	✓
IBM Spectrum Accelerate	✓
IBM Spectrum Scale (ESS and GSS)	
IBM Spectrum Virtualize	✓
IBM Spectrum Virtualize for Public Cloud	✓
SAN Volume Controller	✓
Storwize® V3500	✓
Storwize V3700	✓
Storwize V5000	✓
Storwize V5000E	✓
Storwize V7000	✓
IBMStorwize Flex System V7000 Storage Node	
Storwize V7000 Unified	
XIV®	✓
All others (managed by SMI-S providers)	

Restriction:

- When you add a storage system for monitoring, you must enter a user name and password that is used to connect to that storage system. In rare cases, special characters in the user name or password might not be allowed by IBM Spectrum Control. If you can't add a storage system because of restricted characters, change its credentials and try to add it again. Then, open a ticket for the character limitation so that it can be addressed in a future update.
- IBM Spectrum Control doesn't support monitoring non-IBM software-defined storage devices. However, it can monitor IBM software-defined storage devices, such as IBM SAN Volume Controller. For a list of storage devices that can be monitored, check out <https://www.ibm.com/support/pages/node/6249369>.

Tips:

- To monitor the performance of a Storwize V7000 Unified storage system, you must add it as a block storage system.
- To monitor an IBM Spectrum Virtualize for Public Cloud storage system, you must configure it for communication with IBM Spectrum Control. For more information, see [Configuring IBM Spectrum Virtualize for Public Cloud](#).
- For a complete list of the versions of storage systems and SMI-S providers that you can add, and the firmware levels of supported storage systems, see [IBM Spectrum Control interoperability matrix for storage systems](#).

Procedure

- To add block storage systems, go to Storage > Block Storage Systems. To add a file storage system, go to Storage > File Storage Systems. To add an object storage system, go to Storage > Object Storage Systems.
- Click Add Storage Systems or Add Storage System.
- Click the icon for the type of storage system that you want to add.

Click this icon	To add these storage systems
Storwize Family	IBM® Storwize® V3500, Storwize V3700, Storwize V5000Storwize V7000
SVC or Spectrum Virtualize	SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, or IBM Spectrum Virtualize software-only clusters
FlashSystem Family	FlashSystem 5000 (formerly known as IBM Storwize V5000), , FlashSystem 5200, FlashSystem 7200 (formerly known as IBM Storwize V7000), FlashSystem 7300, FlashSystem 9100, FlashSystem 9200, FlashSystem 9500, FlashSystem V9000, FlashSystem 900, FlashSystem A9000, FlashSystem A9000R
Dell EMC	<ul style="list-style-type: none"> Unity VMAX family with support for SMI-S 1.6 VNX family with support for SMI-S 1.6 VNXe family with support for SMI-S 1.6 All other Dell EMC storage systems with support for SMI-S 1.2
Hitachi	Hitachi VSP
NetApp	<ul style="list-style-type: none"> ONTAP 9 ONTAP 8 (7- mode)
Pure	<ul style="list-style-type: none"> FlashArray//M FlashArray//X
Others	Other non-IBM storage systems

- Specify the connection details and the authentication credentials for the storage systems that you want to add, then click Connect. You can close the window and continue your work. To change the default probes and performance monitors, keep the window open and update them when the storage systems are connected.

Results

The storage systems are added for monitoring. Each storage system is automatically added to the default alert policy for the storage system type. A probe is automatically run to collect status and asset information about the storage systems.

Storage systems can be configured for block storage, file storage, a combination of block and file storage, or a combination of file and object storage. The method that a storage system uses for managing data determines the internal resources and object resources that are monitored.

After you add storage systems, a task is added to the list of running tasks on the page banner. If the task fails, the task is added to the list of failed tasks on the page banner. You can retry failed tasks or clear them from the list.

What to do next

When the collection of data is complete, you can view status information and capacity data about the storage system on the resource list page for the block, file, or object storage system. For example, for block storage systems, you can view the information on the Block Storage Systems page.

After the probe is run, agentless servers are updated with any new information about the host connections on the storage systems. You can review the agentless servers on the Servers page, and make any changes that are required.

Managing alert conditions and notification settings: Review the alert configuration for the storage system. You can change which alert policy manages the storage system. You can also set a storage system to not be managed by any policy.

- [Hitachi](#)
Add Hitachi Virtual Storage Platform (VSP) F and G Series storage systems to get performance, asset, capacity, and configuration metadata analyzed for block storage so that you can detect performance issues, changes in storage usage, and plan for future storage needs.
- [Configuring IBM Spectrum Virtualize for Public Cloud for monitoring](#)
Configure IBM Spectrum Virtualize for Public Cloud for monitoring with IBM Spectrum Control.
- [User roles for collecting performance metadata from IBM Spectrum Virtualize](#)
When you add IBM Spectrum Virtualize storage systems, you must add the credentials of a user who can log on to the storage system and collect performance metadata.
- [Monitoring IBM Spectrum Scale without requiring root privileges](#)
You can enable a non-root user on a IBM Spectrum Scale cluster node to monitor IBM Spectrum Scale storage systems.
- [Verifying that asset, capacity, and configuration metadata can be collected for object storage](#)
To probe object storage, IBM Spectrum Control server must be able to connect to the OpenStack Swift and Keystone endpoints that are used to access object services.
- [Configuring the collection of performance data for IBM Spectrum Scale](#)
Configure the collection of information about the performance of IBM Spectrum Scale file storage systems.
- [Configuring OpenStack access to monitor the object storage system](#)
Configure OpenStack access for the user name that is used to monitor the IBM Spectrum Scale object storage system.

Related information

- [Collecting data](#)

Hitachi

Add Hitachi Virtual Storage Platform (VSP) F and G Series storage systems to get performance, asset, capacity, and configuration metadata analyzed for block storage so that you can detect performance issues, changes in storage usage, and plan for future storage needs.

You can add a VSP system by connecting to the Hitachi Command Suite that is managing the device. By default, the asset, capacity, and configuration metadata that is collected from storage systems is refreshed every 24 hours. Performance metadata is collected at a default collection interval of 5 minutes.

Use the following information to add VSP storage systems so that they can be monitored and their metadata can be collected, analyzed, and presented in the GUI. Add a VSP storage system by connecting to the Hitachi Command Suite that manages the device. Connect to the Hitachi Device Manager to enable the Export Tool to collect performance information from the storage system.

Supported versions: To view the versions of Hitachi VSP storage systems that are supported in IBM Spectrum® Control, go to the [Hitachi support page](#).

Restriction: Performance monitoring is not supported when IBM Spectrum Control is installed on AIX®.

Complete the following steps to add your Hitachi VSP storage systems for monitoring.

- Add Hitachi VSP F and G Series by connecting to Hitachi Command Suite. See [Credentials for Hitachi Command Suite](#).
- Enable performance monitoring while you are adding your Hitachi VSP systems. See [Performance](#).
- Install the Hitachi Export Tool to collect performance information and export it from your storage system. See [Installing the Hitachi Export Tool](#).
- To enable or disable performance monitoring, see [Creating performance monitors in IBM Spectrum Control](#) and [Starting and stopping performance monitors](#).

Add Hitachi VSP F and G Series

When you select Hitachi VSP, you can add one or more storage systems by entering the connection information for the Hitachi Command Suite that is managing those devices.

Important: To discover and monitor Hitachi VSP F and G Series storage systems, you must provide a user name for connecting to the associated Hitachi Command Suite. This user name must be assigned to the View Group user group and have AdminGroup permission in the Hitachi Command Suite.

Credentials for Hitachi Command Suite

IBM Spectrum Control monitors capacity and configuration information about Hitachi devices by using the Hitachi Command Suite. Provide the connection information for the Hitachi Command Suite installation that is monitoring the devices.

Host name or IP address

The host name or IP address that you use to connect to Hitachi Command Suite. Depending on what is supported in your environment, you can enter an Internet Protocol version 4 (IPv4) or IPv6 address. If you enter an IPv6 address, the preferred representation is written as eight groups of four hexadecimal digits. Example: 2001:DB95:0000:1234:0000:0000:5678:ABCD.

User Name and Password

The account credentials that are required to connect to Hitachi Command Suite. The role or user group that is assigned to the user name must have the appropriate privileges to monitor the data that is collected and, if required, to change the frequency of the data collection schedules. [Learn more about the role requirements for the user name.](#)

Performance

Enable performance monitoring

When you enable performance monitoring in IBM Spectrum Control, the user name and password that you provided to connect to the Hitachi Command Suite are also used to connect to Hitachi Device Manager. Hitachi Device Manager provides the connection information for the storage system to the Export Tool, which then sends the performance information to IBM Spectrum Control.

Requirement: You might already have created credentials for Hitachi Device Manager, but you must create new credentials to match the credentials used for Command Suite. Device Manager uses these credentials exclusively for the Export Tool.

Important: To collect performance metadata, Hitachi requires that the password for connecting to the storage system be temporarily stored as clear text in a file on the server where IBM Spectrum Control is installed. For security reasons, ensure access to that server is restricted to key personnel. The file is automatically deleted after metadata is collected.

User Name and Password

The account credentials that are required to connect to Hitachi Device Manager. The role or user group that is assigned to the user name must have the appropriate privileges to monitor the data that is collected and, if required, to change the frequency of the data collection schedules. [Learn more about the role requirements for the user name.](#)

Capacity

Asset, capacity, and configuration metadata is collected by probes. By default, a probe collects metadata from Hitachi VSP systems once every 24 hours.

Configure

If you want to change the default probes and performance monitors, keep the Add Storage System window open. When the storage system or systems are connected, select the systems that you want to configure. To select multiple storage systems, press Ctrl and click or Shift and click. The default configuration settings are applied to the remaining storage systems.

Display Name

The name of the storage system that is displayed in the IBM Spectrum Control interface. If you do not enter a value, the default name is provided. This is the name that is provided with the storage system.

Location

The location, such as the geographical location or the building, where the resource is located.

Data Collection

Schedule the probe and performance monitor for the storage system. The probe collects status, asset, and storage information about the storage system. The performance monitor collects metrics that measure the performance of the storage system.

Probe

Enter the time and schedule for the storage system probe to run.

Enter the time in *hh:mm* format, where *hh* equals the hour and *mm* equals the minute. The time zone that is shown is determined by the location of the Data server for IBM Spectrum Control.

Performance Monitor

If you enabled performance monitoring when you added the Hitachi storage systems, you can change the performance monitoring interval from every 5 minutes to every 60 minutes. You can also disable performance monitoring.

- **[Installing the Hitachi Export Tool](#)**

To collect performance data from a Hitachi storage system, you must install the Hitachi Export Tool on the server or virtual machine where IBM Spectrum Control is installed. If you have many Hitachi storage systems in your environment, you must install the appropriate version of the Export Tool for each model of storage system.

Installing the Hitachi Export Tool

To collect performance data from a Hitachi storage system, you must install the Hitachi Export Tool on the server or virtual machine where IBM Spectrum® Control is installed. If you have many Hitachi storage systems in your environment, you must install the appropriate version of the Export Tool for each model of storage system.

About this task

To collect performance metadata, the Hitachi Export Tool requires that the password for connecting to the storage system is temporarily stored as clear text in a file on the server or virtual machine where IBM Spectrum Control is installed. For security reasons, ensure access to that server or virtual machine is restricted to key personnel. The file is automatically deleted after metadata is collected.

Procedure

1. Use your Hitachi account to download the Export Tool from the following location:
https://knowledge.hitachivantara.com/Knowledge/Storage/How_to_Download_the_Appropriate_Export_Tool_Version_Specific_to_Array_Microcode
Tip: If you have more than one model of Hitachi storage system, you must download the version of the Export Tool for each model.
2. To install the tool into the default directory, first create the following directory structure where *hitachi_model* is the directory name that corresponds to the Export Tool for the model of your storage system.

Windows

Program Files\monitor\export\hitachi_model

Linux®

/opt/monitor/export/hitachi_model

Where *hitachi_model* is the directory name. See the following table for the directory details.

Table 1. Hitachi storage system models and export tool directories

Model	Property Name	Directory Name
Hitachi Virtual Storage Platform 5100/5100H/5500/5500H	hitachi_5000_ExportToolPath	5000
Hitachi Virtual Storage Platform E990	hitachi_E_ExportToolPath	eseries
Hitachi VSP G200/400/600/800 and F400/600/800	hitachi_FGX00_ExportToolPath	fgx00
Hitachi VSP G130/350/370/700/900 and F350/370/700/900	hitachi_FGXX0_ExportToolPath	fgxx0
Hitachi VSP G1000/1500 and F1500	hitachi_FG1X00_ExportToolPath	fg1x00
Hitachi VSP N400/600/800	hitachi_N_ExportToolPath	nseries
Hitachi Virtual Storage Platform (VSP)	hitachi_LegacyVSP_ExportToolPath	legacyvsp

3. Ensure that the Export Tool directory has write permission for users.

4. For each Export Tool that you downloaded, open the disk image and extract the tool for your chosen operating system.

5. Install the tool into the default directory.

6. Optional: If it is not possible to run the Export Tool from the default directory, another directory can be used. However, you must modify the setup.properties file in the DataCollector/conf directory to specify the new location.

Modify the *hitachi_model_ExportToolPath* property for your specific Hitachi storage system as in the following examples for Hitachi VSP G200/400*/600/800 and F400/600/800.

Windows

```
hitachi_FGX00_ExportToolPath=C:\User1\HitachiExportTool\export\fgx00
```

Where folders are separated by \\

Linux

```
hitachi_FGX00_ExportToolPath=/usr/abc/HitachiExportTool/export/fgx00
```

7. Optional: If you changed the setup.properties file, save it and then restart the device server.

Adding and configuring NetApp resources

After adding and configuring a NetApp resource, you can use IBM Spectrum® Control to administer and monitor the status, configuration, capacity, performance, and information that is collected about the resource.

About this task

A NetApp ONTAP 9 resource can be added to IBM Spectrum Control as either a block or file storage system to view all of the block and file data on the system. A NetApp ONTAP 8 (7- mode) resource can be added to IBM Spectrum Control as a block storage device, a file storage device, or both. IBM Spectrum Control collects different information from the device, depending on the type of storage system you specify. For all NetApp storage systems IBM Spectrum Control collects data by directly connecting to the resource.

Important:

- For new IBM Spectrum Control installations, you can no longer deploy Storage Resource agents for monitoring NetApp resources. You must now monitor those resources through an SMI-S provider or a direct connection. However, if you upgrade from a previous version of the product, you can continue to manage NetApp resources that have Storage Resource agents deployed from that previous version.
- When you upgrade from a previous version to IBM Spectrum Control 5.2.10 or later, the existing alert definitions for NetApp filers are migrated to the new version. However, to generate alerts for those filers, you must monitor them through direct connections rather than through Storage Resource agents. To set up direct connections for those NetApp filers, add them again for monitoring and enter the credentials for connecting to them directly. After you add the NetApp filers, their previous entries are automatically removed and alerts will be generated for the alert definitions that were migrated.
- For a complete list of NetApp devices that can be monitored, see [NetApp support page](#).
- For NetApp storage systems managed by SMI-S providers, performance monitoring is supported only for volumes and front-end ports. It is not supported at the storage system level.

To add a NetApp resource for monitoring by IBM Spectrum Control, complete the following steps:

Procedure

- To add a NetApp resource as a block storage system, go to Storage > Block Storage Systems. To add the resource as a file storage system, go to Storage > File Storage Systems.
- Click Add Storage System.
- Click the NetApp icon to add the resource as a storage system.
- Complete the connection information for the storage system and schedule the collection of data.

Results

After the initial discovery has completed, you can view information about the NetApp device in the File Storage Systems or Block Storage Systems pages in the GUI.

Antivirus software restriction: If your antivirus software is set on the maximum mode, it might prevent your ONTAP 9 storage systems from being added to IBM Spectrum Control. For more information about how to configure your antivirus software, see [Installation checklists for IBM Spectrum Control](#).

What to do next

By default, asset, capacity, and configuration metadata is aggregated and collected daily. You can schedule daily capacity and inventory reports to gain insights about your NetApp storage systems.

Configuring IBM Spectrum Virtualize for Public Cloud for monitoring

Configure IBM Spectrum Virtualize for Public Cloud for monitoring with IBM Spectrum® Control.

About this task

IBM Spectrum Virtualize is a software-defined storage solution that has been proven for years in SAN Volume Controller and the IBM® Storwize® family. IBM Spectrum Virtualize for Public Cloud extends that solution to a hybrid-cloud or cloud-based model, where servers, storage, and network infrastructure are delivered in a public cloud environment. It can be deployed on either IBM® Cloud or Amazon Web Services (AWS) cloud infrastructures.

With IBM Spectrum Control, you can view the capacity, space usage, and performance of your IBM Spectrum Virtualize for Public Cloud storage systems. Other monitoring features, such as alerting, health checking, advanced analytics, and reporting are also supported.

Before you can add an IBM Spectrum Virtualize for Public Cloud storage system for monitoring, you must ensure that IBM Spectrum Control can connect to it. To enable a connection, you can use the site-to-site VPN IPsec tunnel that exists between the on-premises environment and the IBM Spectrum Virtualize for Public Cloud storage systems.

- **Monitoring IBM Spectrum Virtualize for Public Cloud with on-premises data collection (Site to Site VPN IPsec)**
You can connect and monitor IBM Spectrum Virtualize for Public Cloud storage with IBM Spectrum Control by using the site-to-site virtual private network (VPN) IPsec tunnel that exists between the on-premises environment and the IBM Spectrum Virtualize for Public Cloud instances.

Monitoring IBM Spectrum Virtualize for Public Cloud with on-premises data collection (Site to Site VPN IPsec)

You can connect and monitor IBM Spectrum Virtualize for Public Cloud storage with IBM Spectrum® Control by using the site-to-site virtual private network (VPN) IPsec tunnel that exists between the on-premises environment and the IBM Spectrum Virtualize for Public Cloud instances.

About this task

The VPN IPsec site-to-site tunnel is a secure communication network between the cloud infrastructure and the on-premises environment. Network communication between the private subnets is controlled by the access control list (ACL) that is populated when you create the VPN IPsec site-to-site tunnel.

Typically, a bi-directional, IPsec site-to-site tunnel is limited to the subnets that contain the following IP addresses:

- On-premises IBM Spectrum Virtualize cluster and replication target
- Cloud-based IBM Spectrum Virtualize for Public Cloud cluster and replication target

To use the IPsec site-to-site tunnel for communication between IBM Spectrum Control and IBM Spectrum Virtualize for Public Cloud, you must include the IP addresses of your IBM Spectrum Control servers (or the subnet in which they reside) in the tunnel definition as one of the on-premises endpoints.

For example, in the AWS Management Console, you can define the following:

- The external (internet-routable) IP address of the on-premises IPsec tunnel endpoint.
- The subnet / IP addresses for the on-premises IBM Spectrum Control servers, IBM Spectrum Virtualize cluster, and replication target that will communicate through the tunnel.
- The IP addresses of the IBM Spectrum Virtualize for Public Cloud cluster and replication target in AWS that will communicate through the tunnel.

When defined, you can export a configuration file that can be used in different IPsec VPN devices, such as Vyatta and Juniper (Junos VPN Site Secure). Contact your network and firewall administrators to help set up this communication between IBM Spectrum Control and IBM Spectrum Virtualize for Public Cloud.

Procedure

After configuring the connection between IBM Spectrum Virtualize for Public Cloud and IBM Spectrum Control, add the storage system for monitoring by completing the following steps:

1. In IBM Spectrum Control, go to **Storage > Block Storage Systems**.
2. Click **Add Storage Systems**.
3. Click the SVC or Spectrum Virtualize icon.
4. Specify the IP address and authentication credentials for the IBM Spectrum Virtualize for Public Cloud instance that you want to monitor.
5. Click **Connect**.

Results

The storage system is added for monitoring and is automatically added to the default alert policy for the storage system type. Data collection is automatically run to collect status, configuration, capacity, and performance metadata about the storage system.

What to do next

By default, asset, capacity, and configuration metadata is aggregated and collected daily. Performance metadata is collected every 5 minutes. You can schedule daily capacity and inventory reports to gain insights about your IBM Spectrum Virtualize for Public Cloud storage systems.

User roles for collecting performance metadata from IBM Spectrum Virtualize

When you add IBM Spectrum Virtualize storage systems, you must add the credentials of a user who can log on to the storage system and collect performance metadata.

The required role of the user on your storage system for collecting performance metadata depends on the version of IBM Spectrum Virtualize.

Versions earlier than 8.3.1.2

The user must have the role of Administrator or SecurityAdmin.

Version 8.3.1.2 or later

The user can have any role, such as the Monitor role.

Important information about user roles for IBM Spectrum Virtualize 8.3.1.2 or later

If the collection of performance metadata is stopped on the storage system, the role of the storage system user determines whether the collection can be automatically restarted. For example, if the storage system user has the Monitor role and your administrator has stopped the collection on the storage system, the collection must be manually restarted by your administrator before IBM Spectrum® Control can collect performance metadata.

User has privileges to start collections

If the user has the Administrator or SecurityAdmin role on the storage system, the collection is automatically restarted and the performance metadata is collected.

User does not have privileges to start collections

If the user has the Monitor role on the storage system, IBM Spectrum Control can't automatically restart the collection. To manually start the collection of performance metadata on the storage system, complete the following steps:

1. Log on to the storage system as a user with the Administrator, or SecurityAdmin role.
2. Run the following command:

```
svctask startstats -interval time in minutes
```

The interval must be less than or equal to the performance monitor interval for the storage system in IBM Spectrum Control. To check the value, in the IBM Spectrum Control GUI, go to the Block Storage Systems page and view the value for Performance Monitor Interval for the storage system.

Monitoring IBM Spectrum Scale without requiring root privileges

You can enable a non-root user on a IBM Spectrum Scale cluster node to monitor IBM Spectrum Scale storage systems.

About this task

Before you add the IBM Spectrum Scale storage system for monitoring, complete this task.

To grant access to the user, make the following changes to the sudoers file on the cluster node that is used for authentication:

- Add the set of administration commands that are required to monitor the storage system.
- Associate the user with the set of administration commands.

Important:

- You must configure IBM Spectrum Scale cluster nodes for SSH login without requiring a password.
- You can't monitor IBM Spectrum Scale clusters that are configured with a sudo wrapper environment. This configuration is not supported because mmdsh* commands are required to collect configuration information from the IBM Spectrum Scale cluster.

Procedure

1. Log on to the cluster node that is used for authentication with a user name that has root privileges.
2. To edit the sudoers file, enter the following command:

```
visudo -f /etc/sudoers
```

3. Add the following command aliases to the `sudoers` file.

The command aliases contain the commands that the user must be able to issue to monitor the storage system.

Important: Each command alias must be on a single line without line breaks.

```
Cmdnd Alias TPC_GPFS_MMCD = /usr/lpp/mmfs/bin/mmsdrquery, /usr/lpp/mmfs/bin/mmfsconfig,  
/usr/lpp/mmfs/bin/mmgetstate, /usr/lpp/mmfs/bin/mmfsnodeclass, /usr/lpp/mmfs/bin/mmfsfs,  
/usr/lpp/mmfs/bin/mmdf, /usr/lpp/mmfs/bin/mmfsnsd, /usr/lpp/mmfs/bin/mmfsfileset,  
/usr/lpp/mmfs/bin/mmcloudgateway, /usr/lpp/mmfs/bin/mmfsmount, /usr/lpp/mmfs/bin/mmfsnapshot,  
/usr/lpp/mmfs/bin/mmrepquota, /usr/lpp/mmfs/bin/mmfspolicy, /usr/lpp/mmfs/bin/mmapplypolicy
```

```
Cmdnd Alias TPC_GPFS_MMDSH = /usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmdiag --version,  
/usr/lpp/mmfs/bin/mmdsh -N * /lib/udev/scsi_id --whitelisted *,  
/usr/lpp/mmfs/bin/mmdsh -N * /sbin/blockdev --getsize64 *,  
/usr/lpp/mmfs/bin/mmdsh -N * /usr/bin/getconf DISK_SIZE *,  
/usr/lpp/mmfs/bin/mmdsh -f 20000 -N linuxNodes 'cat /sys/class/fc_host/*',  
/usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces node list,  
/usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces service list -a,  
/usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces address list|grep object_database_node,  
/usr/lpp/mmfs/bin/mmdsh -N * /usr/lpp/mmfs/bin/mmces address list  
--by-node|grep object_database_node,  
/usr/lpp/mmfs/bin/mmdsh -v -N cesNodes /usr/lpp/mmfs/bin/mmobj config list  
--ccrfile object-server.conf --section DEFAULT --property devices,
```

```

/usr/lpp/mmfs/bin/mmdsh -f 20000 -v -N * "test -e /opt/IBM/zimon/ZIMonSensors.cfg &&
(grep -w collectors -A 4 /opt/IBM/zimon/ZIMonSensors.cfg | grep -w host) || true",
/usr/lpp/mmfs/bin/mmdsh -f 20000 -v -N nonWindowsNodes hostname

Cmd Alias TPC_GPFS_MMDSH2 = /usr/lpp/mmfs/bin/mmdsh -f 20000 -v -N localhost test -e /opt/IBM/zimon/ZIMonSensors.cfg &&
(grep \
 \ -w 'collectors' -A 4 /opt/IBM/zimon/ZIMonSensors.cfg \ | grep -w 'host') || true

Cmd Alias TPC_GPFS_MMDSH3 = /usr/lpp/mmfs/bin/mmdsh -f 20000 -v -N linuxNodes test -e /opt/IBM/zimon/ZIMonSensors.cfg &&
(grep \
 \ -w 'collectors' -A 4 /opt/IBM/zimon/ZIMonSensors.cfg \ | grep -w 'host') || true

Cmd Alias TPC_GPFS_OTHER = /bin/cat *release,/usr/bin/lsb_release -a, /bin/date, /usr/bin/date, /bin/grep,
/bin/true, /usr/bin/test

Cmd Alias TPC_GPFS_CMDS = TPC_GPFS_MMCMDS, TPC_GPFS_MMDSH, TPC_GPFS_OTHER, TPC_GPFS_MMDSH2, TPC_GPFS_MMDSH3

```

4. To enable the user to issue the commands, add the following lines after the command aliases:

```

Defaults:user_name !requiretty
user_name ALL=(ALL) TPC_GPFS_CMDS

```

where *user_name* is the user name that you enter when you add the storage system for monitoring.

Results

The user that you added to the sudoers file can now monitor the IBM Spectrum Scale storage system.

Verifying that asset, capacity, and configuration metadata can be collected for object storage

To probe object storage, IBM Spectrum® Control server must be able to connect to the OpenStack Swift and Keystone endpoints that are used to access object services.

Procedure

To verify that the server can access the object services, complete these steps:

1. List the URLs for the Keystone and Swift services by running commands on a IBM Spectrum Scale cluster node that is configured for object storage. Run the following commands with a user name that has root privileges:

```

. ~/openrc
openstack endpoint list

```
2. Ensure that the IBM Spectrum Control server can connect to the IP addresses and host names that are included in the Keystone and Swift services URLs. For example, the URL for the Keystone service might be `http://gpfs420proto1:5000/v3`. The IBM Spectrum Control server must be able to connect to the `gpfs420proto1` host name.

Related concepts

- [Collecting asset and status data by using IBM Spectrum Control probes](#)

Configuring the collection of performance data for IBM Spectrum Scale

Configure the collection of information about the performance of IBM Spectrum Scale file storage systems.

About this task

To collect performance data, you must first configure the IBM Spectrum Scale performance monitoring tool on the IBM Spectrum Scale cluster. Then, use the IBM Spectrum® Control GUI to schedule the collection of performance data.

Procedure

1. Configure the IBM Spectrum Scale performance monitoring tool by using the **mmperfmon config** command. For more information on the configuration code list, see the documentation for the https://www.ibm.com/docs/en/STXKQY_5.1.1/com.ibm.spectrum.scale.v5r10.doc/bl1adm_mmperfmon.htm. Set the **--collector** property of **mmperfmon config** to the IBM Spectrum Scale cluster node where the collector component will run. The **--collectors** property must be set to one of the following options:
 - An IP address that can be reached by the IBM Spectrum Control server
 - A host name that resolves to an IP address that is reachable by the IBM Spectrum Control server

Ensuring that the collector node can be reached by IBM Spectrum Control:
If the **--collectors** property is set to the internal host name of the collector node, the collector node might not be reachable by the IBM Spectrum Control server.

Update the hosts file on the IBM Spectrum Control server to resolve the internal host name of the collector node to the public IP address that the IBM Spectrum Control server can access.

2. Configure the collector component of the performance monitoring tool so that the IBM Spectrum Control server can connect to port 9084 on the cluster node where the collector component will run.
Configure the collector component by editing the `/opt/IBM/zimon/ZIMonCollector.cfg` file on the IBM Spectrum Scale cluster node where the collector component will run.

IBM Spectrum Scale 5.1.0 or earlier

To collect performance metadata, ensure that the `queryinterface` property is set to "0.0.0.0" in `ZIMonCollector.cfg` file on the cluster node:

```
queryinterface = "0.0.0.0"
```

IBM Spectrum Scale 5.1.1 or later

To collect performance metadata, ensure that `zimon/ZIMonCollector.cfg` file on the cluster node includes the following properties:

```
fallbackqueryinterface = "0.0.0.0" # "0.0.0.0" to allow remote connections (or ":::0" for IPv6)
fallbackqueryport = "9084"
```

3. Restart the collector component.
4. Use the `mmperfmon` command to assign the `perfmon` designation to each node for which you want to collect performance data. For more information, see the documentation for the <https://www.ibm.com/docs/en/spectrum-scale/5.1.1?topic=reference-mmchnode-command>
Note: By default in IBM Spectrum Scale 5.0.0 or higher, the sensors are started on all nodes.
5. Schedule the collection of performance data by creating a performance monitor in the IBM Spectrum Control GUI.
You can create performance monitors in the GUI when you add resources for monitoring or you can create them later.

Results

A performance monitor is created for the IBM Spectrum Scale storage system. If a successful probe run is completed for the resource, the performance monitor runs according to the defined interval.

What to do next

To check the progress of a performance monitor, you can complete the following actions:

- View the Performance Monitor Status column on the File Storage Systems page.
- From the File Storage Systems page, right-click a row, and select Data Collection > Open Performance Monitor Logs.

To learn more about configuring the performance monitoring tool, see https://www.ibm.com/support/knowledgecenter/STXKQY_5.0.1/com.ibm.spectrum.scale.v5r01.doc/bl1adv_PMToverview.htm.

Related tasks

- [Collecting performance data by using IBM Spectrum Control performance monitors](#)

Related reference

- [Performance metrics](#)

Configuring OpenStack access to monitor the object storage system

Configure OpenStack access for the user name that is used to monitor the IBM Spectrum Scale object storage system.

About this task

Before you add the IBM Spectrum Scale object storage system for monitoring, ensure that the user name that you use to monitor the storage system has access to object storage accounts.

Tip: The terms "account" and "project" mean the same thing. Swift, the OpenStack object storage service, uses the term "account" and Keystone, the OpenStack identity service, uses the term "project."

Procedure

1. Set the object storage account and the domain for the user. By default, the domain is set to Default and cannot be modified after you create the user account.
Choose one of the following actions:

- a. To set the account and domain when you create a user account, use the following command:

```
openstack user create --domain domainname --project projectname
--password Password username
```

- b. To set the account and the domain for an existing user name, use the following command:

```
openstack user set --project projectname username
```

2. Assign the admin role for an object storage account to the user name.
Use the following command:

```
openstack role add --user username --project projectname admin
```

3. To monitor all accounts on the object storage system, assign to the user name the role that is defined in the `reseller_admin_role` configuration option in the Swift proxy server. The default value for the `reseller_admin_role` option is `ResellerAdmin`.

For example, use the following command:

```
openstack role add --user username --project projectname ResellerAdmin
```

Restriction: If you do not assign the ResellerAdmin role, information is collected only for the object storage accounts that the user has admin access to.

What to do next

To learn more about configuring and assigning user roles in OpenStack, see <https://docs.openstack.org/>.

Adding fabrics and switches

Add fabrics and switches for monitoring.

Before you begin

When you add switches for monitoring, you must provide a user name and password for logging in to those switches. The role or user group of the user determines the type of data that can be collected. To see the privileges that are required for switch users, see [Required user roles for monitoring resources](#).

For a Cisco switch to successfully receive and respond to queries from IBM Spectrum® Control, the following basic requirements must be met:

- IBM Spectrum Control can use SNMPv3 (preferred) or SNMPv1 to probe switches and fabrics. The SNMPv3 protocol is preferred because it provides better security, but switches that use the SNMPv1 protocol are also supported. Some switches are configured to use SNMPv3 by default.
- The Fibre Alliance FC Management MIB (FA MIB) and Fibre Channel Fabric Element MIB (FE MIB) must be enabled on the switch.
- When you use the SNMPv1 protocol, the community string that is configured in IBM Spectrum Control must match one of the community strings that are configured on the switch with read access. Additionally, Cisco switches must have a community string match for write access. The default community strings in IBM Spectrum Control are "public" for read access and "private" for write access. Other community strings can be defined on the switches, but are not used. Community strings are not relevant when you use the SNMPv3 protocol.
- SNMP access control lists must include the IBM Spectrum Control system. These access control lists are defined and configured on the switches. Some lists automatically include all hosts, while others exclude all by default.
- The Fibre Channel (FC) or Fibre Channel over Ethernet (FCoE) protocols must be enabled on the switch. Some switches, such as the Cisco Nexus 5000 series, require you to enable these protocols. Otherwise, IBM Spectrum Control does not recognize the switch when you try to add it using the Add Switches and Fabrics for Monitoring dialog. For instructions on how to configure Cisco switches for FCoE enablement, go to the Cisco product website at <http://www.cisco.com> and click Support.

IBM Spectrum Control can monitor Brocade switches and fabrics in either of the following ways:

- Use the REST API to manage the switches and fabrics.
- Use the Brocade Network Advisor (BNA) Storage Management Initiative (SMI) agent to manage the switches and fabrics. The SMI agent is embedded in the BNA. BNA can manage multiple fabrics within and across data centers. When you configure BNA, you set up one switch to be the *seed* switch that interconnects to all the other switches in the fabric.

Important: The embedded SMI agent is only available in the Professional Plus and Enterprise editions of BNA. BNA is no longer available for purchase, although it will be supported by Broadcom until February 2022. For more information about end of support for BNA, see <https://www.broadcom.com/support/fibre-channel-networking/eol>.

About this task

In IBM Spectrum Control, you can add switches and fabrics for monitoring at the same time. You can add the following types of switches:

- Brocade
- Cisco

For a complete list of the firmware levels of supported switches, see <https://www.ibm.com/support/pages/node/6249365>.

Procedure

1. From the Network menu, click Switches or Fabrics.
2. Click Add Switch or Add Fabric.
3. Select the vendor: Brocade or Cisco.
4. Specify how you want to discover the switches and the fabrics that are connected to the switches. For Brocade switches and fabrics, you can connect using either REST API or an SMI-S provider on Brocade Network Advisor (BNA). For Cisco switches and fabrics, you must use SNMP agents.

Type of switch	Steps
Brocade	<p>Complete the following steps:</p> <ol style="list-style-type: none">a. Select the version of the Fabric OS that is running on the switches.b. If you select 8.2.1 or later, specify the connection details and the authentication credentials for the switches that you want to add. You need to add only one switch from each fabric to IBM Spectrum Control. If other switches that are on the same fabric are running Fabric OS 8.2.1 or later and have the same username, password, protocol, and port, they are added to IBM Spectrum Control automatically. The user must have a user or admin account with chassis-role privileges for the switches that are added. <p>Click Next.</p> <ol style="list-style-type: none">c. If you select Earlier than 8.2.1, make sure that the switch that you want to add is monitored by BNA. BNA is required for monitoring Brocade switches with Fabric OS versions earlier than 8.2.1. Click Next, then specify the connection details and the authentication credentials for the BNA SMI agent that you want to add. You can close the window and continue your work. To change the default probes and performance monitors, keep the window open and update them when the switches are connected.

Type of switch	Steps
Cisco	<p>To enable the collection of data with SNMP agents, complete the following steps:</p> <ol style="list-style-type: none"> Specify the version of SNMP: v3 (preferred) or v1. If you choose to use v3, specify the User name and associated Authentication password. Select an Authentication protocol and an Encryption protocol. If you choose to use v1, define the Read community and the Write community. Identify the Host name or IP address of each switch that you want to add. Click Next. You can close the window and continue your work. To change the default probes and performance monitors, keep the window open and update them when the switches are connected.

5. Specify the display name and location of the switches and fabrics that are discovered.

Results

The discovered switches are automatically added for monitoring, if they have the same credentials as the switches that you added. A probe is automatically run to collect status and asset information about the resources.

When Brocade switches are added, other switches in the same fabric that do not have the same authentication credentials might also be discovered. These switches are not added for monitoring. They are added to the Switches page with a Not Monitored condition.

What to do next

When the collection of data is complete, you can view status information about the fabric or switch on the Fabrics page or Switches page.

If you want to monitor multiple switches that have Not Monitored status, you must add them with the correct credentials. To add a list of Not Monitored switches, add the IP Address column to the Switches page. Then, select Condition as the filter, and Not Monitored as the filter value. Review the switches and gather the IP addresses of the switches that you want to add for monitoring. Then, add the switches in the usual way.

Managing alert conditions and notification settings: Review the alert configuration for the fabric or switch.

If you configured data collection for the switches or fabric, they are automatically added to the default alert policy for their type.

You can change which alert policy manages a switch, and you can also set a switch to not be managed by any policy. For example, to change the alert policy for a switch, go to Settings > Alert Policies and click the Resources by Policy tab. Click Switches, select a switch, and then select Actions > Set Policy.

Restriction: If you don't configure data collection for the switches and fabric, they are not added to the default alert policy.

- [Data sources for switches and fabrics](#)

Depending on the type of switch that you want to monitor, you can use Fabric OS REST API, SMI agents, or SNMP agents to monitor switches and fabrics. SMI agents are also known as CIM agents, CIMOM (CIM Object Manager) agents, or SMI-S providers.

- [Configuring Brocade switches for monitoring](#)

Configure Brocade switches for monitoring through a direct connection. IBM Spectrum Control establishes this connection by using the Brocade REST API and requires Fabric OS 8.2.1 or later.

Related tasks

- [Sorting lists](#)

Data sources for switches and fabrics

Depending on the type of switch that you want to monitor, you can use Fabric OS REST API, SMI agents, or SNMP agents to monitor switches and fabrics. SMI agents are also known as CIM agents, CIMOM (CIM Object Manager) agents, or SMI-S providers.

The vendors of switches can help you determine which agents to use.

Table 1. Agent types for switch and fabric functions

Function	Brocade	Cisco
Monitor performance	REST API or SMI agent	SNMP agent
Collect information about switches and switch ports	REST API or SMI agent	SNMP agent
Collect information about topology connectivity	REST API or SMI agent	SNMP agent
Collect information about zoning information and zone control	REST API or SMI agent	SNMP agent
Generate alerts	REST API or SMI agent	SNMP agent

Important: The embedded SMI agent is only available in the Professional Plus and Enterprise editions of Brocade Network Advisor (BNA). BNA is no longer available for purchase, although it will be supported by Broadcom until February 2022. For more information about end of support for BNA, see <https://www.broadcom.com/support/fibre-channel-networking/eol>.

Configuring Brocade switches for monitoring

Configure Brocade switches for monitoring through a direct connection. IBM Spectrum® Control establishes this connection by using the Brocade REST API and requires Fabric OS 8.2.1 or later.

About this task

Before you add a Brocade switch for monitoring in IBM Spectrum Control, you must configure it for HTTPS and set up the proper user role.

Procedure

1. To enhance security, configure a switch for HTTPS. This action disables HTTP access.

Option 1

Create a self-signed HTTPS certificate by using the **seccertmgmt generate** command.

```
seccertmgmt generate -cert [https | extn -keypair_tag keypair_tag] [-type [rsa | dsa | ecdsa]] [-keysize value] [-hash type] [-years value] [-f]
```

For example:

```
switch:admin> seccertmgmt generate -cert https
```

Option 2

Upload a CA-signed certificate by using the **seccertmgmt import** command.

```
seccertmgmt import -cert [fcap | commoncert | https | radius | ldap | syslog | extn -keypair_tag keypair_tag | mgmtip] -protocol [scp | ftp] -ipaddr IP address -remotedir remote_directory -certname certificate_name -cacert preimported_local_ca_cert -login login_name -password password
```

For more information, see [Brocade® Fabric OS® Command Reference Manual, 8.2.x](#).

2. Configure a user with required roles for the switches.
 - a. To collect metadata about a switch, IBM Spectrum Control requires a user account with "user" or "admin" role. If the switch is virtualized, the user must also have a "user" or "admin" role for the chassis and have access to all the Logical Fabric IDs 1 - 128.
 - b. To create a user with the required roles, run one of the following commands.
 - For virtualized switches, run the **userconfig --add** command with the **-r role**, **-l LF_ID_LIST**, and **-c chassis_role** options.

```
virtualizedswitch:admin> userconfig --add user_name -r user -l 1-128 -c user -p MyPassword
```
 - For non-virtualized switches, run the **userconfig --add** command with the **-r role** option.

```
physicalswitch:admin> userconfig --add user_name -r user -p MyPassword
```
 - c. To change the role for an existing user, run one of the following commands.
 - For virtualized switches, run the **userconfig --change** command with the **-r**, **-l**, or **-c** options along with the **userconfig --addlf** command to expand the list of Logical Fabric IDs.

For example, run the following command to change the chassis role and the list of Logical Fabric IDs that the user is allowed to access.

```
virtualizedswitch:admin> userconfig --change user_name -c admin -l 128
```

```
virtualizedswitch:admin> userconfig --addlf user_name -c admin -l 1-128
```
 - For non-virtualized switches, run the **userconfig --change** command with the **-r** option.

```
physicalswitch:admin> userconfig --change user_name -r admin
```
 - d. To verify the roles, run the **userconfig --show** command. Add the **-a** option to list all users on the switch.

Adding hypervisors

You can add hypervisors for monitoring by IBM Spectrum Control. Hypervisors include VMware vSphere data sources such as ESX and ESXi, and vCenter Server systems. When you add vSphere data sources, you can then collect data, generate reports, and manage storage that is related to those resources. To add a hypervisor, you specify connection information for the ESX or ESXi data source. You can add multiple hypervisors by specifying connection information for a vCenter Server system.

- [Adding ESX and ESXi hypervisors](#)
Add ESX and ESXi hypervisors for monitoring by IBM Spectrum Control.
- [Adding vCenter Server systems](#)
For IBM Spectrum Control to monitor multiple hypervisors, you must first add a vCenter Server system.

Adding ESX and ESXi hypervisors

Add ESX and ESXi hypervisors for monitoring by IBM Spectrum Control.

About this task

You can add an ESX or ESXi hypervisor by specifying connection information for that hypervisor. For a complete list of hypervisors that you can add, see the [IBM Spectrum Control interoperability matrix](#), and go to the *Agents, Servers and Browsers* section.

Procedure

1. In the menu bar, go to Servers > Hypervisors.
2. Click Add Hypervisor.
3. Select ESX/ESXi and enter connection information about the hypervisor.
4. Schedule a probe of the hypervisor.
5. Follow the instructions in the GUI to add the hypervisor.

Results

The hypervisor is added for monitoring. A probe is automatically run to collect status and asset information about the hypervisor.

What to do next

When the collection of data is complete, you can view status information and capacity data about the hypervisor on the Hypervisors page.

Managing alert conditions and notification settings: Review the alert configuration for the hypervisor. Each hypervisor is automatically added to the default alert policy for the hypervisor type. You can change which alert policy manages a hypervisor, and you can also set a hypervisor to not be managed by any policy. For example, to change the alert policy for a hypervisor, go to Settings, > Alert Policies and click the Resources by Policy tab. Click Hypervisors, select a hypervisor, and then select Actions, > Set Policy.

- [Checking permissions to browse data stores](#)

Determine if the user name that you specified for a VMware data source has permission to browse through the data stores on a hypervisor.

Related tasks

- [Adding vCenter Server systems](#)

Checking permissions to browse data stores

Determine if the user name that you specified for a VMware data source has permission to browse through the data stores on a hypervisor.

About this task

When you add a VMware data source in IBM Spectrum® Control, the user name that you specify must have permission to browse through the data stores on VMware. IBM Spectrum Control must browse through the data stores to collect information from the hypervisors. However, the "Read Only" role as defined by VMware does not allow IBM Spectrum Control to browse the data stores. You can use the "Virtual Machine Power® User" role if you do not want to use the Administrator role, or you can create a custom role with the required permissions.

Procedure

To verify that a VMware user is assigned the correct role and privileges to monitor VMware data sources, follow these steps:

1. Ensure that the user role has the required VMware datastore permissions by completing the following steps:
 - a. Connect the vSphere Web Client to the VMware data source.
The data source can be an ESX server, a vCenter Server Appliance, or a vCenter Server.
 - b. From the Inventories view, select Hosts and Clusters.
 - c. Select a host, and click the Related Objects tab.
 - d. View the datastores by clicking the Datastores tab.
 - e. Right-click a datastore, and select File Browser. If you can view the Files tab for the datastore, your browse permission is working correctly.
2. Determine the role that is assigned to the user by logging in to the vSphere Web Client by using the administrator user ID. From the Administration view, select Roles. Verify the role name that is assigned to the user.
3. Determine the privileges that are assigned to the role by selecting the user's role and clicking Privileges. Expand the privilege groups to view the specific privileges.
4. Optional: If you must edit the privileges for the role, select the role and click the Edit role action icon. Select privilege groups or expand to select specific privileges.

What to do next

For more information about VMware user roles, go to the [VMware documentation center](#) and search for *vSphere users and permissions*.

Adding vCenter Server systems

For IBM Spectrum Control to monitor multiple hypervisors, you must first add a vCenter Server system.

About this task

You can add multiple hypervisors by specifying connection information for a vCenter Server system. For a complete list of hypervisors and vCenter Server systems that you can add, see the [IBM Spectrum Control interoperability matrix](#), and go to the *Agents, Servers, and Browsers* section.

Procedure

1. In the menu bar, go to Servers, > Hypervisors.
2. Click Add Hypervisor.
3. Select VMware vCenter and enter connection information about the vCenter Server system.
4. Schedule a probe for the hypervisors that were discovered.
5. Follow the instructions in the wizard to add the vCenter Server system.

Results

After a vCenter Server system is added for monitoring, probes collect status and asset information about its hypervisors.

What to do next

When the collection of data is complete, you can view status information and capacity data about the hypervisors on the Hypervisors page.

Managing alert conditions and notification settings: Review the alert configuration for the hypervisor.

Each hypervisor is automatically added to the default alert policy for the hypervisor type. You can change which alert policy manages a hypervisor, and you can also set a hypervisor to not be managed by any policy.

You can also create an alert policy from the alert definitions and notification settings of a hypervisor. Click Servers > Hypervisors. Right-click the hypervisor from which you want to create an alert policy, then click View Alert Definitions. Click Create Policy from the Policy Actions menu.

Related tasks

- [Checking permissions to browse data stores](#)

Adding servers

Add servers for monitoring. IBM Spectrum® Control creates and updates agentless servers automatically after it probes storage systems and hypervisors. You can also add a server by deploying a Storage Resource agent.

- **Agentless servers**
IBM Spectrum Control creates and updates agentless servers automatically after it probes storage systems and hypervisors.
- **Adding servers with Storage Resource agents**
You can deploy a Storage Resource agent to a server if you want to view information about the resources that are internal to a server, or about the fabrics that are visible to the server. Internal resources include controllers, disks, file systems, logical volumes, shares, and directories.

Agentless servers

IBM Spectrum® Control creates and updates agentless servers automatically after it probes storage systems and hypervisors.

IBM Spectrum Control creates agentless servers by using information about the following resources:

- Host connections on storage systems
- Host bus adapters (HBAs) on fabrics
- Discovered virtual machines

Depending on your storage environment, IBM Spectrum Control might not model all of the servers in your environment correctly. For example, IBM Spectrum Control might identify six host connections and create an agentless server for each host connection. However, in your environment, these six host connections represent one server computer. In this case, you need to merge the six agentless servers into one agentless server.

Similarly, IBM Spectrum Control might create a single agentless server from a group of host connections that appear to be related. However, in your environment, each of these host connections represents a separate server computer. In this case, you can separate the agentless server into multiple agentless servers, one for each of the servers in your environment.

Adding servers with Storage Resource agents

You can deploy a Storage Resource agent to a server if you want to view information about the resources that are internal to a server, or about the fabrics that are visible to the server. Internal resources include controllers, disks, file systems, logical volumes, shares, and directories.

Before you begin

To deploy a Storage Resource agent, you must have Administrator privileges.

For each installation of IBM Spectrum® Control, you can deploy only one Storage Resource agent on each server. If you attempt to deploy additional Storage Resource agents on a server, the deployments fail.

Procedure

1. In the menu bar, go to Servers > Servers.
2. Click Deploy Agent.
3. Select one of the following methods for deploying Storage Resource agents:
 - Enter information about the server and the Storage Resource agent manually.
 - Import configuration information from a comma-delimited file.
4. On the Deploy Agent page, configure deployment information for the Storage Resource agents.
If you deploy multiple agents with different operating systems, separate configuration pages are displayed for agents that are deployed on Windows servers and agents that are deployed on UNIX servers.
5. On the Configure page, schedule the deployment of the Storage Resource agents.
If you are deploying agents on multiple servers, a time span is calculated during which the agents are deployed. The agents are deployed at regular intervals during the time span to avoid excessive load on the IBM Spectrum Control server.

6. Schedule the time and frequency that probes are run for the servers.
If you add multiple servers, a time span is calculated during which the servers are probed.
7. To add the servers, click Finish.

Results

A probe is automatically run for a server after the agent is successfully deployed.

What to do next

Use Agent State on the Servers page to monitor the status of the agent deployment.

Managing alert conditions and notification settings: Review the alert configuration for the server.

Each server is automatically added to the default alert policy for the server type. You can change which alert policy manages a server, and you can also set a server to not be managed by any policy.

You can also create an alert policy from the alert definitions and notification settings of a server. Click Servers > Servers. Right-click the server from which you want to create an alert policy, then click View Alert Definitions. Click Create Policy from the Policy Actions menu.

- [File List](#)
Add one or more servers with Storage Resource agents by importing the configuration information from a comma-delimited file.

File List

Add one or more servers with Storage Resource agents by importing the configuration information from a comma-delimited file.

The web-based GUI guides you through the following steps for adding servers:

- Select the input file.
- Configure deployment information.
- Schedule the agent deployment and data collection for the servers.

The comma-delimited file that you use to import the configuration information for servers can contain entries for a single server or multiple servers. Each line in the file represents a server that you want to add. The information about each server must be organized in the following format:

```
host name or IP address,OS type,location,custom tag 1,custom tag 2,custom tag 3
```

where:

- **host name or IP address** is required for each server entry. An IP address can be in an IPv4 or IPv6 format. A host name or IP address can contain the following characters or symbols:
 - A - Z (uppercase characters)
 - a - z (lowercase characters)
 - 0 - 9 (numeric characters)
 - Symbols: - . : _
- **OS type** is required and represents the operating system of the server. The OS type for a server must be one of the following values:
 - Windows
 - Linux®
 - AIX®
- **location** is optional and represents the physical location of the server. The location value can be up to 64 characters in length. If the length exceeds 64 characters, the location value is truncated when the server is added.
- **custom tag 1, custom tag 2, and custom tag 3** are optional and represent any additional information that you want to provide about the server. The custom tag values can be up to 64 characters in length. If the length exceeds 64 characters, the custom tag value is truncated when the server is added.
Tip: The custom tags can be displayed on the Servers page or can be included as report columns when you generate reports for the server.

If a server has no information for an optional field, you must add a comma separator unless there are no further values for that server.

Example

```
host1,Windows,San Jose,Accounting department
host5,Linux,London,Finance department
198.51.100.22,AIX,,Computing department
2001:DB8:0:0:0:0:0:0,Windows,Tokyo
```

Tips:

- If the comma-delimited file contains entries for multiple servers, it might take some time to add the servers. To confirm that the servers are added, check the Status column on the Servers page.
- To comment out a line, enter a "#" at the beginning of the line. The server on that line is not added when the list is imported.
Example:

```
# host1,Windows,San Jose,Accounting department
```

- If there are syntax problems in the file, none of the servers in the file are added.

Adding rollup servers

Rollup servers gather capacity and status information from IBM Spectrum® Control servers in your enterprise. You can add rollup servers for monitoring on the Settings > Rollup Server Connections page.

IBM Spectrum Control supports primary and secondary rollup servers. Primary rollup servers gather capacity and status information about the storage resources that are managed by secondary rollup servers. You can view this information by enabling IBM Spectrum Control rollup mode. These managed resources include servers, hypervisors, fabrics, switches, and block storage systems.


Configuring rollup servers

To configure rollup servers, you add one or more secondary rollup servers to a primary rollup server. The primary rollup server uses probes to collect capacity and status information from the secondary rollup servers. Primary rollup servers must have IBM Spectrum Control 5.3.0 or later installed and running. Secondary rollup servers must have IBM Spectrum Control 5.3.0 or later installed.

Adding a rollup server

To add and configure secondary rollup servers, go to Settings > Rollup Server Connections. You can add one or more secondary rollup servers to the primary rollup (local) server. You can also add a primary rollup server as a secondary rollup server to another primary server.

Viewing rollup information

To view information about the resources managed by your rollup servers, you must be in *rollup mode*. To enter rollup mode, hover your pointer over the rollup mode icon  on the menu bar and select Enter rollup mode. You cannot enter rollup mode unless you have added at least one secondary rollup server to the primary rollup (local) server.

For more information about rollup servers and rollup mode, see [Rollup servers](#).

- **Restrictions for rollup servers**

When you use rollup mode to view information that was collected by secondary servers and rolled up to primary servers, keep in mind the following limitations, restrictions, and considerations.

Restrictions for rollup servers

When you use rollup mode to view information that was collected by secondary servers and rolled up to primary servers, keep in mind the following limitations, restrictions, and considerations.

Required versions of IBM Spectrum® Control

Primary rollup servers must have IBM Spectrum Control 5.3.0 or later installed. You can add secondary rollup servers that run any version of IBM Spectrum Control 5.3.0 or later.

Collecting data after you upgrade

After you upgrade a primary rollup server you must run probes of the secondary rollup servers that were added to the primary server prior to the upgrade. Use the Start Probe action to run probes of the secondary rollup servers. Wait for the first probe (manual or scheduled) to complete, before you work with a new function.

Unmapped Capacity and Mapped Capacity values are slightly different between the secondary and primary rollup servers

Due to differences in the methods that are used to calculate Unmapped Capacity and Mapped Capacity, there might be minor variations between the primary and secondary rollup servers. The value that is recorded in the secondary rollup server resource list panel is the definitive one.

Removing resources

Remove resources that you no longer want to monitor with IBM Spectrum® Control.

Procedure

1. In the menu bar, go to the type of resource that you want to remove.

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics
Rollup subordinate servers	Settings > Rollup Server Connections

2. Right-click the resource and click Remove.

Tips:

- To remove rollup subordinate servers, click Actions > Remove Server.
- To remove a Cisco fabric, you must remove all the switches in that fabric. The Cisco fabric is then automatically removed.
- To remove a NetApp storage system configured for block storage, or configured for block storage and file storage, go to Storage > Block Storage Systems. If the NetApp storage system is configured only as a filer, then go to Storage > File Storage Systems.

3. Click Remove to confirm that you want to remove the resource.

Results

When you remove a resource, it is no longer monitored by IBM Spectrum Control. All the data that is associated with the resource, including historical data, is removed from the database. All the scheduled tasks and data collection jobs for the resource are canceled.

Collecting data

Determining the data that you want to gather about storage resources is critical to helping you implement a storage management strategy. IBM Spectrum® Control provides two different jobs for collecting data about resources.

Table 1. Data collection jobs, resources, and related tasks

Data collection job	Resources	Related tasks
Probe Use probes to collect status, asset, configuration, and capacity information about resources.	Storage systems Servers Hypervisors Fabrics Switches Rollup servers (capacity and status information only)	<ul style="list-style-type: none">Schedule a probe when you add a resource for monitoring.Modify the schedule for a probe at any time on the list and details pages for a resource. For example, to modify the schedule for a block storage system, go to the Block Storage Systems page or the details page for the block storage system.Use the Rollup server connections page to manage probes for rollup servers. Learn more...
Performance monitor Use performance monitors to collect metrics that measure the performance of resources.	Block storage systems File storage systems	<ul style="list-style-type: none">Schedule a performance monitor when you add a resource for monitoring.Modify the schedule for a performance monitor at any time on the Performance page, and on the list and details pages for the resource. Learn more...

- [Collecting asset and status data by using IBM Spectrum Control probes](#)
Use IBM Spectrum Control probes to collect asset, status, and storage data about monitored resources in your environment. Probes are also used to discover information about new or removed disks and file systems. You can view information about probes on the resource pages in IBM Spectrum Control, for example, the Block Storage Systems page, or the Block Storage Systems details page.
- [Collecting performance data by using IBM Spectrum Control performance monitors](#)
Before you can monitor the performance of the resources in your environment, you must collect data about those resources. IBM Spectrum Control uses performance monitors to collect metrics for measuring the performance of storage systems and switches. Use the Performance Monitors page to view and manage performance monitors.
- [Collecting information about shares on storage systems](#)
To collect information about shares in file systems, you must modify the TPCD.config file. When you provision shares, you can then view such information as the name of the share, the name of the file system, and the protocols that are used to share files.
- [Collecting information about the sizes of snapshots in IBM Spectrum Scale](#)
By default, IBM Spectrum Control does not collect information about the size of GPFS snapshots in IBM Spectrum Scale. To view that information, you must enable IBM Spectrum Control to collect information about the size of snapshots.

Related tasks

- [Viewing capacity alerts and violations](#)
- [Defining alert definitions for general attributes and capacity changes](#)

Collecting asset and status data by using IBM Spectrum Control probes

Use IBM Spectrum® Control probes to collect asset, status, and storage data about monitored resources in your environment. Probes are also used to discover information about new or removed disks and file systems. You can view information about probes on the resource pages in IBM Spectrum Control, for example, the Block Storage Systems page, or the Block Storage Systems details page.

- [Creating probes in IBM Spectrum Control](#)
Use probes to collect asset, status, and storage data about resources that are monitored by IBM Spectrum Control. You can create probes when you add resources for monitoring or after you add the resources for monitoring.
- [Verifying that a probe is running for a resource](#)
Verify that a probe is running for a resource to ensure that IBM Spectrum Control is collecting asset and status information about the resource.
- [Modifying probes](#)
Probes are data collection jobs that collect status and asset information about monitored resources in your environment. To modify a probe in IBM Spectrum Control, use the resource list page for the monitored resource in the web-based GUI.
- [Configuring alerts for probes](#)
Configure the alert notifications that are generated if a probe fails to run.
- [Starting probes](#)
Use the Start Probe action to immediately collect data about resources.
- [Viewing probes for a specific resource](#)
You can view information for specific resources on the resource details pages. For example, use the Storage Systems details page to view the date and time when data was last collected.
- [Viewing probe logs](#)
A probe log file is created for each run of a probe on a storage resource. Probe logs include detailed information about the status, actions, and progress of a probe.

Related tasks

- [Verifying that asset, capacity, and configuration metadata can be collected for object storage](#)

Creating probes in IBM Spectrum Control

Use probes to collect asset, status, and storage data about resources that are monitored by IBM Spectrum® Control. You can create probes when you add resources for monitoring or after you add the resources for monitoring.

About this task

You can create probes for one or more resources at the same time. You can also create an automated or manual probe. Schedule an automated probe to manage your probe jobs efficiently and to minimize the load on the IBM Spectrum Control server. If you want to run a probe at a specific time, schedule a manual probe.

Procedure

1. To schedule a probe for a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics
Rollup subordinate servers	Settings > Rollup server connections

Tip: For switches, probes also collect asset and topology information about the fabrics that include those switches.

2. Locate the resources that you want to probe.
3. Select the resource rows and click Actions > Data Collection > Schedule. For Rollup subordinate servers, select Actions > Start probe for the server.
4. On the Data Collection Schedule window, set the probe to Enabled.
5. From the following table, choose the type of probe that you want. Make your selections for the data collection, and save your selections.

Type of probe	Selections for probe
Automated probe	Select the frequency of the data collection. Click Save.
Manual probe	Select the time and the frequency of the data collection. Click Save.

For automated probes, the start time is determined by an algorithm that minimizes the number of concurrent probe jobs to avoid excessive load on the IBM Spectrum Control server. Automated probe jobs are scheduled by IBM Spectrum Control within the time interval that you define in the Automated Probe Run Window.

Tips for manually scheduling probes:

- When you schedule probes to collect data about resources, schedule the probes to run at different times. By scheduling data collection so that no two probes run at the same time, you can help improve the overall performance of the product and ensure that data is collected successfully.
- If you schedule manual probes for multiple resources, a time span is calculated during which the resources are probed. The resources are probed at different times during the time span to avoid excessive load on the IBM Spectrum Control server.

Results

Probe schedules are created for the resources that are selected. The probe jobs run according to the modified schedule. If you disabled the probes, IBM Spectrum Control stops collecting asset, status, and capacity data for the monitored resources.

Note: Status and asset information about logical switches is collected when the parent switch is probed.

Example

You select three storage systems, `storage_system_a`, `storage_system_b`, and `storage_system_c`. A probe schedule was defined for `storage_system_c`, but not for `storage_system_a` and `storage_system_b`. You set the time and frequency fields, and set the probe to Enabled. When you click Save, probe schedules are created for `storage_system_a` and `storage_system_b`. The probe schedule that was defined for `storage_system_c` remains unchanged.

• How automated probes are scheduled

When you schedule probes, you can manually specify when they run, or you can have IBM Spectrum Control automatically schedule them. Automated probe scheduling helps remove the complexity of manually planning when probes are run and ensures that server resources (CPU, Memory, and so on) required for data collection are distributed over time. Distributing the use of system resources over time helps avoid excessive load on IBM Spectrum Control at any one time.

How automated probes are scheduled

When you schedule probes, you can manually specify when they run, or you can have IBM Spectrum® Control automatically schedule them. Automated probe scheduling helps remove the complexity of manually planning when probes are run and ensures that server resources (CPU, Memory, and so on) required for data collection are distributed over time. Distributing the use of system resources over time helps avoid excessive load on IBM Spectrum Control at any one time.

To schedule probes automatically, an algorithm is used to determine the best times for when the probes are run. To determine this time, the algorithm considers the existing schedules of probes for all types of resources. For example, when an automated probe of a storage system is scheduled, IBM Spectrum Control attempts to minimize the number of probes that are running concurrently by analyzing the existing probes of switches, servers, hypervisors, and other storage systems.

IBM Spectrum Control also analyzes the frequency, start time, and expected duration of existing probe schedules. This analysis includes probes that were scheduled automatically or manually. After the analysis, IBM Spectrum Control creates a unique probe schedule for each day of the week.

Keep in mind the following information about automatic probe scheduling:

- When you schedule a new probe during the week, the schedule of existing probes does not change. Existing probes are considered when the schedule of the new probe is determined, but they continue to run according to the previously defined schedule for that week.
- When a probe is scheduled to run on a frequency of every X days where $X > 1$, IBM Spectrum Control determines the days for that frequency based on a number of conditions. Some of the conditions include whether the probe was run previously as a scheduled probe and how much time is available until the end of the current probe window.
- Automated probes can be rescheduled by IBM Spectrum Control under the following conditions:
 - When you change the time range for the automated probe schedule, IBM Spectrum Control reschedules all the automated probes to optimize when probes are run during that new time range.
 - On Saturday night (in the time zone of the IBM Spectrum Control server) of every week, all automated probes are rescheduled. During the rescheduling process, the algorithm is again applied to determine the best times to run automated probes on each day of the upcoming week. Because all automated probes are rescheduled weekly, IBM Spectrum Control ensures that the average number of concurrent probes is kept minimized for your environment on an ongoing basis.

Verifying that a probe is running for a resource

Verify that a probe is running for a resource to ensure that IBM Spectrum® Control is collecting asset and status information about the resource.

About this task

To verify that a probe is running for a resource, use the resource list page in the web-based GUI.

Procedure

1. From the menu bar, go to a resource list page for a resource type.
For example, to verify that a probe is running for a storage system, go to Storage and select the type of storage system you want to monitor.
2. Locate the resource that you want to check.
3. View the Probe Status column.
A value of **Successful** or **Running** indicates that the probe is collecting asset and status information about the resource. The Last Successful Probe column shows the most recent date and time when the probe successfully collected data.
4. Optional: If the Probe Status column shows a status of **Failed** or **Warning**, right-click the resource and click Data Collection > Open Probe Logs to view the messages in the log file.
5. Optional: If the Probe Status column shows a status of **Never probed**, right-click the resource and click Data Collection > Schedule to schedule a probe for the resource.

Modifying probes

Probes are data collection jobs that collect status and asset information about monitored resources in your environment. To modify a probe in IBM Spectrum® Control, use the resource list page for the monitored resource in the web-based GUI.

About this task

You can select one or more devices and modify their schedules.

Procedure

1. To modify a probe for a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Object storage systems	Storage > Object Storage Systems
Servers	Servers > Servers
Hypervisors	Servers > Hypervisors
Switches	Network > Switches
Fabrics	Network > Fabrics

Tip: For switches, probes also collect asset and topology information about the fabrics that include those switches.

2. Locate the resources that you want to probe.
3. Select the resource rows and click Actions > Data Collection > Schedule.
4. On the Data Collection Schedule window, modify the schedule for the probes, and click Save.

For automated probes, the start time is determined by an algorithm that minimizes the number of concurrent probe jobs to avoid excessive load on the IBM Spectrum Control server. Automated probe jobs are scheduled by IBM Spectrum Control within the time interval that you defined in the Automated Probe Run Window.

If you schedule manual probes for multiple resources, a time span is calculated during which the resources are probed. The resources are probed at different times during the time span to avoid excessive load on the IBM Spectrum Control server.

The following rules determine the settings for the probe fields when you open the Data Collection Schedule window:

- The Enable Probe check box is checked by default. Uncheck this box to disable the probe operation for all selected resources.
- If all selected resources are configured with automated probe scheduling, then Automatically is selected; if all selected resources are configured with manual probe scheduling, then Manually is selected. If some resources are configured with automatic probe scheduling and some resources are configured with manual probe scheduling, then neither Automatically nor Manually is selected.

- Fields that are blank indicate that more than one value (Automatically or Manually) is used among the selected resources. If the field contains a red asterisk (*), you must set a value. If the field is blank, the existing values are preserved unless you select a new value to be applied to all of the selected resources.

Results

The probe jobs run according to the modified schedule. If you disabled the probes, IBM Spectrum Control stops collecting asset and status data for the monitored resources.

The changes that you make to the probe schedule are applied to all the resources that you select. For example, you select two servers with different probe frequency values. Probes are disabled for both servers. When you open the Data Collection Schedule window, the probe status field is set to Disabled and the frequency field is blank. You set the status field to Enabled and the frequency field to **Every 2 days**. When you click Save, the probe is enabled with the new frequency value for the two servers that you selected

Configuring alerts for probes

Configure the alert notifications that are generated if a probe fails to run.

About this task

You can define alerts for a resource if you want to be notified when the following conditions occur:



- The status of its probe is error or warning. An error status occurs when a probe did not complete and no data was collected about a resource. A warning status occurs when a probe completes, but might not have collected a complete set of data.
- A specified amount of time has passed since a probe collected data about a resource. You can use this type of alert to be notified when up-to-date configuration and status data is not being collected about a resource and its existing data might be stale.

Procedure

To define an alert for the probe that collects data about a resource, complete the following steps:

- To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage_> Block Storage Systems
File storage systems	Storage_> File Storage Systems
Object storage systems	Storage_> Object Storage Systems
Servers	Servers_> Servers
Hypervisors	Servers_> Hypervisors
Switches	Network_> Switches
Fabrics	Network_> Fabrics

- Right-click a resource in the list and click Edit Alert Definitions.
- To enable the Probe Status alert for an attribute, click the corresponding switch icon  and select a status to be alerted on.
You can use this alert to be notified when a probe is only collecting partial data about a resource or not collecting any data.
- To enable the Last Successful Probe alert for an attribute, click the corresponding switch icon  and specify how much time can pass before an alert is generated.
You can use this alert to be notified when up-to-date configuration and status data is not being collected about a resource and its existing data might be stale.
- Click Save.

Starting probes

Use the Start Probe action to immediately collect data about resources.

Before you begin

If a probe job is not defined for the resource, use the Schedule action to create a probe.

About this task

You can start probes for multiple resources at the same time.

Tip: When you start a probe for a switch, the probe gathers statistics about the fabric that the switch is a part of.

Procedure

- From the menu bar, go to the resource list page for a resource type.
For example, to start a probe for a storage system, go to Storage and select the type of storage system you want to monitor.
- Locate the resources that you want to probe.
- Select the resource rows and click Actions_>Data Collection_>Start Probe.
The Start Probe action is available for the resources that you select if at least one of the resources meets the following criteria:
 - A probe job is defined for the resource.
 - The probe job is enabled.
 - The probe job is not currently running.

Results

The probe status changes to Running.

Probes are started for all the resources that meet the criteria. For example, you select two storage systems, **Storage_System_a** and **Storage_System_b**. Probes are enabled for both of the storage systems and a probe is running for **Storage_System_a**. The probe is started for **Storage_System_b** when you select Start Probe. **Storage_System_a** is not affected by the action because a probe is already running for the resource.

Viewing probes for a specific resource


You can view information for specific resources on the resource details pages. For example, use the Storage Systems details page to view the date and time when data was last collected.

Procedure

1. In the menu bar in the web-based GUI, go to a type of top-level resource.
For example, if you want to view the resource information for a storage system, go to Storage and select the type of storage system you want to monitor.
2. Locate the resource that you want to view information for.
3. Right-click the resource row and select View Details.
4. In the General section of the resource details page, click Data Collection.

Results




The Data Collection pane is displayed, which contains information about the data collection jobs that are defined for the resource.

Note: To view information about rollup resources, you must be in *rollup mode*. To enter rollup mode, move the mouse pointer over the rollup icon  in the menu bar and select Enter rollup mode.

Viewing probe logs

A probe log file is created for each run of a probe on a storage resource. Probe logs include detailed information about the status, actions, and progress of a probe.

Procedure

1. In the menu bar in the web-based GUI, go to a resource list page for a resource type.
For example, to view the probe logs for a storage system, go to Storage and select the type of storage system you want to monitor.
2. Locate the resource that you want to view.
3. Right-click the resource row, and select Data Collection > Open Probe Logs.
Restriction: If you select multiple rows, the Open Probe Logs action is not available.
The Logs page opens and shows log entries for the most recent run of the probe.
4. Optional: On the Logs page, to view a previous probe log, select a log from the Select a log list.
5. Optional: To view only the log entries that have a Warning or Error status, select an option from the Show all list.
You can choose to view only entries that have the following statuses:
 -  Only error entries
 -  Only warning entries
 -  Error and warning entries
6. Optional: To view an explanation of the message that is associated with a log entry, click the link in the ID column.

What to do next

You can view the following information on the Logs page:

- The overall status of a probe. The icon that is shown in the Select a log list represents the most critical status that was generated by an action in the job run.
- The status for each probe action.
- The date and time when an action was completed. The date, time, and time zone of an action is shown in the Date and Time column.
- The description of an action.

Collecting performance data by using IBM Spectrum Control performance monitors

Before you can monitor the performance of the resources in your environment, you must collect data about those resources. IBM Spectrum® Control uses performance monitors to collect metrics for measuring the performance of storage systems and switches. Use the Performance Monitors page to view and manage performance monitors.

About this task

You must schedule and run performance monitors before you can complete other tasks, such as optimizing storage tiering, balancing pools, and running performance reports.

Before you can view performance information for resources, you must complete the following tasks:

- Add the storage system or switch for monitoring by IBM Spectrum Control.
- Schedule a performance monitor to collect performance data about the resource.
- Define a performance alert so that you are notified when the performance of a resource might represent a potential problem. For example, you can define an alert threshold that notifies you when the total I/O port rate for a storage system falls outside a specified range.

The performance information that is gathered includes metrics that measure the performance of volumes, ports, and disks. You can view performance metrics on resource list and detail pages for switches and storage systems. For example, you can view performance metrics that are related to switches in the Performance tab on the Switches page.

- [Creating performance monitors in IBM Spectrum Control](#)
Performance monitors are data collection jobs that collect metrics for measuring the performance of storage systems and switches. You can create performance monitors in IBM Spectrum Control when you add resources for monitoring or at a later time.
- [Verifying that a performance monitor is running for a resource](#)
Verify that a performance monitor is running for a resource to ensure that IBM Spectrum Control is collecting performance metrics about the resource. Use performance monitors to collect metrics about the performance of storage systems and switches.
- [Modifying performance monitors](#)
To modify all of the performance monitors that are used by IBM Spectrum Control, use the Performance Monitors page in the web-based GUI. To modify performance monitors for specific resources, use the resource list and resource detail pages.
- [Configuring alerts for performance monitors](#)
Configure the alert notifications that are generated if a performance monitor fails to run.
- [Starting and stopping performance monitors](#)
To start or stop performance monitors for storage systems and switches, use the Performance Monitors page in the GUI.
- [Viewing performance monitors](#)
To view all of the performance monitors that IBM Spectrum Control uses to collect metrics for measuring the performance of resources, use the Performance Monitors page. To view the performance monitors for specific resources, use resource list and resource details pages.

Related tasks

- [Configuring the collection of performance data for IBM Spectrum Scale](#)
- [Defining alert definitions for performance changes](#)

Related reference

- [Performance metrics](#)

Creating performance monitors in IBM Spectrum Control

Performance monitors are data collection jobs that collect metrics for measuring the performance of storage systems and switches. You can create performance monitors in IBM Spectrum® Control when you add resources for monitoring or at a later time.

Before you begin

To create performance monitors, you must have Administrator privileges.

About this task

To create performance monitors for resources in your environment, you can use the list and detail pages for the resources. For example, to create performance monitors for switches, use the Switches page. You can create performance monitors for multiple resources at the same time.

Restrictions:

- You can define only one performance monitor for a resource.
- You can create performance monitors for the following resources only:
 - Storage systems
 - Switches that are managed by a CIM agent

Procedure

1. From the menu bar, go to a resource list page for a resource type.
For example, to create performance monitors for storage systems, go to Storage and select the type of storage system you want to monitor.
2. Locate the resources that you want to monitor.
3. Select the resource rows and click Actions, > Data Collection > Schedule.
4. On the Data Collection Schedule window, set the performance monitors to Enabled.
You can set a performance monitor to Enabled only if the probe is set to Enabled.
5. Select the interval for the performance monitors, and click Save.
The interval represents the number of minutes over which samples of performance data are averaged. If you are creating performance monitors for multiple resources, the interval list displays the performance monitor intervals that are common to all resources.
For example, you select two storage systems. **Storage_System_a** specifies performance monitor intervals of 5 minutes, 10 minutes, and 15 minutes. **Storage_System_b** specifies intervals of 10 minutes, 15 minutes, and 20 minutes. The interval list displays intervals of 10 minutes and 15 minutes.
Restrictions:
Note the following restrictions when you create performance monitors:
 - a. If you specify a 1-minute interval for performance monitors, data collection samples might be inconsistent lengths. The inconsistent lengths occur because it might take longer than 1 minute to collect the data on some systems. The inconsistent lengths can cause the data in the historical line chart to be

misleading. This problem can also occur for the 5-minute summary chart when a 1-minute interval was specified for data collection. Check the messages in the job log to determine the length of time that data collection takes for a resource. If the data collection takes longer than 1 minute, consider changing the interval to 5 minutes. You can use a 1-minute interval when you want to troubleshoot a resource.

- b. If you specify a 1-minute interval, the amount of data that is stored in the database repository increases significantly. IBM Spectrum Control stores only 7 days of sample data that is collected at 1-minute intervals.

Results

Performance monitors are created for the resources that are selected, and that meet the criteria. For example, you select three storage systems, **Storage_System_a**, **Storage_System_b**, and **Storage_System_c**. You enable the performance monitors and set the interval field. When you click Save, performance monitors are created for **Storage_System_a** and **Storage_System_b**. **Storage_System_c** is not affected by the action because you cannot configure performance monitors for that storage system.

If a successful probe run is completed for a resource, the performance monitor runs according to the defined interval.

What to do next

To check the progress of a performance monitor, you can complete the following actions:

- From the resource list page, right-click a resource row, and select **Data Collection > Open Performance Monitor Logs**. You can view detailed informational, warning, and error messages that are related to the performance monitor. Use this information to troubleshoot any errors that might occur when the performance monitor runs.
- View the Performance Monitor Status column on the resource list page.

Verifying that a performance monitor is running for a resource

Verify that a performance monitor is running for a resource to ensure that IBM Spectrum® Control is collecting performance metrics about the resource. Use performance monitors to collect metrics about the performance of storage systems and switches.

About this task

To verify that a performance monitor is running for a resource, use the resource list page in the web-based GUI.

Procedure

1. From the menu bar, go to a resource list page for a resource type.
For example, to verify that a performance monitor is running for a storage system, go to **Storage** and select the type of storage system you want to monitor.
2. Locate the resource that you want to check.
3. View the Performance Monitor Status column.
A value of **Completed** or **Running** indicates that the performance monitor is collecting performance metrics about the resource. The **Last Successful Monitor** column shows the most recent date and time when the performance monitor successfully collected data.
4. Optional: If the Performance Monitor Status column shows a status of **Running with problems**, **Completed with warnings**, **Failed**, or **Canceled**, you can view the logs to troubleshoot the error or warning messages. To view the performance monitor logs, right-click the resource row and click **Data Collection > Open Performance Monitor Logs**.
5. Optional: If the Performance Monitor Status column shows a status of **Disabled**, to enable the performance monitor for the resource, right-click the resource row and click **Data Collection > Schedule**. On the **Data Collection Schedule** page, select **Enabled**.
Tip: You can set the performance monitor to **Enabled** only if the probe is set to **Enabled**. If a probe is set to **Disabled**, the performance monitor cannot run.

Related tasks

- [Defining application alerts for performance metrics](#)
- [Defining general group alerts for performance metrics](#)

Modifying performance monitors

To modify all of the performance monitors that are used by IBM Spectrum® Control, use the **Performance Monitors** page in the web-based GUI. To modify performance monitors for specific resources, use the resource list and resource detail pages.

Before you begin

To modify performance monitors, you must have Administrator privileges.

About this task

You can modify the schedule for multiple performance monitors at the same time.

You can enable or disable performance monitors and modify the interval. The interval represents the number of minutes over which samples of performance data are averaged.

Procedure

1. In the menu bar in the web-based GUI, go to Home > Performance Monitors.
Detailed information about all performance monitors is shown in the Performance Monitors tab.
 2. Locate the performance monitors that you want to modify.
 3. Select the performance monitor rows and click Actions > Schedule.
 4. On the Data Collection Schedule window, modify the schedule for the performance monitors.
If you select multiple performance monitors, the following rules are used to determine the field settings when you open the Data Collection Schedule window:
 - If all of the performance monitors that you select are enabled, the Enabled value is displayed. If all of the performance monitors are disabled, the Disabled value is displayed. Otherwise, the state field for the performance monitors is blank.
 - If the performance monitors have different interval values, the interval field is blank. If the performance monitors have the same interval value, the value is displayed.Intervals that are common to all performance monitors that you select are displayed in the interval list. For example, you select performance monitors for two storage systems. **Storage_System_a** specifies performance monitor intervals of 5 minutes, 10 minutes, and 15 minutes. **Storage_System_b** specifies intervals of 10 minutes, 15 minutes, and 20 minutes. The interval list displays intervals of 10 minutes and 15 minutes.
5. Click Save.

Results

The changes are applied to all the performance monitors that are selected. The performance monitors run according to the modified schedule. If you disabled a performance monitor, IBM Spectrum Control stops collecting performance metrics for the storage system or switch.

What to do next

- To check the progress of the performance monitors, you can complete the following actions on the Performance Monitors page:
- View information such as the status of the performance monitor and the most recent date and time when performance data was collected about a resource.
 - Right-click the performance monitor row, and select Open Logs. You can view detailed informational, warning, and error messages that are related to the performance monitor job. Use this information to troubleshoot any errors that might occur when the job runs.

Configuring alerts for performance monitors



Configure the alert notifications that are generated if a performance monitor fails to run.

About this task

- You can define alerts for a storage system or switch if you want to be notified when the following conditions occur:
- The status of its performance monitor is error or warning. An error status occurs when a performance monitor did not complete and no performance data was collected about a resource. A warning status occurs when a performance monitor completes, but might not have collected a complete set of performance data.
 - A specified amount of time has passed since a performance monitor collected data about a resource. You can use this type of alert to be notified when up-to-date performance data is not being collected about a resource and its existing data might be stale.

Procedure

- To define an alert for the probe that collects data about a resource, complete the following steps:
1. To select a resource, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems
Switches	Network > Switches
 2. Right-click a resource in the list and click Edit Alert Definitions.
 3. To enable the Performance Monitor Status alert for an attribute, click the corresponding switch icon  and select a status to be alerted on.
You can use this alert to be notified when a performance monitor is only collecting partial data about a resource or not collecting any data.
 4. To enable the Last Successful Monitor alert for an attribute, click the corresponding switch icon  and specify how much time can pass before an alert is generated.
You can use this alert to be notified when up-to-date performance data is not being collected about a resource and its existing data might be stale.
 5. Click Save.

Starting and stopping performance monitors

To start or stop performance monitors for storage systems and switches, use the Performance Monitors page in the GUI.

Before you begin

To start or stop performance monitors, you must have Administrator privileges.

About this task

You can start or stop performance monitors for multiple resources at the same time.

Restriction: A successful probe must be completed on the resource before a performance monitor can start. If a successful probe is not run, you cannot start the performance monitor for the resource.

Procedure

1. In the menu bar in the GUI, go to Home > Performance Monitors.
2. Click the Performance Monitors tab.
3. Locate the performance monitors that you want to start or stop.
4. Select the performance monitor rows and click Actions > Start or Actions > Stop.

Results

Performance monitors are started or stopped for all the resources that meet the criteria. For example, you select two storage systems, **Storage_System_a** and **Storage_System_b**. Performance monitors are enabled for the storage systems and both resources are probed. A performance monitor is running for **Storage_System_a** but not for **Storage_System_b**. The performance monitor is started for **Storage_System_b** when you select Start. **Storage_System_a** is not affected by the action because a performance monitor is already running for the resource.

Starting or stopping a performance monitor has the following effects:

Start

Starting a performance monitor starts the immediate collection of data from the monitored resource. The status of the performance monitor changes to Running.

Stop

Stopping a performance monitor stops the collection of performance data from the resource. If a performance monitor is running, it stops. Performance data is not collected from the resource until you restart the performance monitor. The status of the performance monitor changes to Canceled.

What to do next

To check the progress of the performance monitor, you can complete the following actions on the Performance Monitors page:

- View information such as the status of the performance monitor and the percentage of data collections that succeeded during the most recent 24 hours when the performance monitor was active.
- Right-click the performance monitor row, and select Open Logs. You can view detailed informational, warning, and error messages that are related to the performance monitor. Use this information to troubleshoot any errors that might occur when the performance monitor runs.

Viewing performance monitors

To view all of the performance monitors that IBM Spectrum® Control uses to collect metrics for measuring the performance of resources, use the Performance Monitors page. To view the performance monitors for specific resources, use resource list and resource details pages.

- [Viewing all performance monitors](#)
To view all of the performance monitors that IBM Spectrum Control uses to collect metrics for measuring the performance of storage systems and switches, use the Performance Monitors page.
- [Viewing performance monitors for specific resources](#)
To view the performance monitors for specific storage systems and switches, use the resource detail pages in the web-based GUI. For example, use the Storage Systems detail page to view performance monitor information for a specific storage system.
- [Viewing performance monitors logs](#)
Use performance monitor logs to view detailed information about the status, actions, and progress of a performance monitor. You can use this information to troubleshoot any errors that might occur when a performance monitor is running.

Viewing all performance monitors

To view all of the performance monitors that IBM Spectrum® Control uses to collect metrics for measuring the performance of storage systems and switches, use the Performance Monitors page.

About this task

The Performance Monitors page displays one row for each storage system or switch that IBM Spectrum Control is monitoring. If a performance monitor is not scheduled for a resource, the status column is set to Disabled for the performance monitor row.

Before IBM Spectrum Control can collect performance metrics for a resource, the resource must fulfill the following conditions:

- A probe must be run for the resource.
- For switches, the switch must be managed by a CIM agent.

Procedure

1. In the menu bar, go to Home > Performance Monitors.
On the Performance Monitors tab, the information about performance monitors is organized into columns. These columns include the status of performance monitors, the names of the monitored resources, and the most recent dates and times when data was successfully collected.
2. Optional: View the status icons at the top of the page to view a summary of performance monitor statuses.
This summary includes the number of performance monitors that have a Running, Running with problems, Failed, or Not Running status.

Viewing performance monitors for specific resources

To view the performance monitors for specific storage systems and switches, use the resource detail pages in the web-based GUI. For example, use the Storage Systems detail page to view performance monitor information for a specific storage system.


About this task

If performance monitor information is not displayed on the resource detail page, this indicates that the resource does not support the collection of performance metrics by IBM Spectrum® Control.

Procedure

1. In the menu bar in the web-based GUI, go to a type of top-level resource.
For example, if you want to view the performance monitor for a storage system, go to Storage and select the type of storage system you want to monitor.
2. Right-click a resource in the list, and select View Details.
The resource detail page is displayed.
In the General section, the number next to Data Collection represents the number of data collection jobs that are associated with the resource. For example, Data Collection (2) indicates that you can view a performance monitor and a probe for the resource.

The icon next to the Data Collection link represents the most critical status of the data collection jobs that are associated with the resource. For example, if a job failed, the following icon is shown:

Data Collection (2) 
3. In the General section, click Data Collection to view information about the performance monitor for the switch or storage system.
On the Data Collection pane, you can view information such as the status of the most recent run of a performance monitor, the performance monitor interval, and the most recent data and time that data was successfully collected.
Tip: You can use the columns on the resource list pages to view information about the performance monitors that are defined for storage systems or switches. For example, use the Storage Systems page to view performance monitor information for storage systems.





Viewing performance monitors logs

Use performance monitor logs to view detailed information about the status, actions, and progress of a performance monitor. You can use this information to troubleshoot any errors that might occur when a performance monitor is running.

About this task

Performance monitor logs contain informational, warning, and error messages that are related to each action that is taken during the processing of a performance monitor.

Procedure

1. In the menu bar in the web-based GUI, go to Home > Performance Monitors.
2. In the Performance Monitors tab, locate the performance monitor that you want to view.
3. Right-click the performance monitor row, and select Open Logs.
Restriction: If you select multiple rows, the Open Logs action is not available.
The Logs page opens and shows log entries for the most recent performance monitor log. A new log file is created when you stop and restart a performance monitor.
4. Optional: On the Logs page, to view a previous performance monitor log, select a log from the Select a log list.
5. Optional: To view only the log entries that have a Warning or Error status, select an option from the Show all list.
You can choose to view only entries that have the following statuses:
 -  Only error entries
 -  Only warning entries
 -   Error and warning entries
6. Optional: To view an explanation of the message that is associated with a log entry, click the link in the ID column.

Results

The information on the Logs page is automatically updated every 30 seconds. New entries are added to the end of a log. You can view the following information on the Logs page:

- The overall status of a performance monitor. The icon that is shown in the Select a log list represents the most critical status that was generated by an action in the job run.
- The status for each performance monitor action.
- The date and time when an action was completed. The date, time, and time zone of an action is shown in the Date and Time column.
- The ID of the message that is associated with an action. You can click the value in the ID column to view more information about a message.
- The description of an action.

Tip: You can use the navigation and search functions of the web browser to locate information on the page.

Collecting information about shares on storage systems

To collect information about shares in file systems, you must modify the TPCD.config file. When you provision shares, you can then view such information as the name of the share, the name of the file system, and the protocols that are used to share files.

About this task

Set the `saveNonRoot` parameter to `saveNonRoot=1` in the TPCD.config file to view the information that is collected about shares in the GUI.

Procedure

1. Open the TPCD.config file.

For Windows operating systems:

By default, the TPCD.config file is in the `installation_dir\data\config\` directory.

For Linux® or UNIX operating systems:

By default, the TPCD.config file is in the `installation_dir/data/config/` directory.

Where `installation_dir` is where IBM Spectrum® Control is installed.

2. In the **[server]** section of the TPCD.config file, change the value of the `saveNonRoot` parameter to 1 as shown in the following example:

```
saveNonRoot=1
```

3. Save the TPCD.config file.
4. Restart the Data server.

Results

When you provision new shares, you can view information about the shares on the Provisioning task page in the GUI. On the Shares page, you can view more information about shares, for example, the servers that have access to the share.

Collecting information about the sizes of snapshots in IBM Spectrum Scale

By default, IBM Spectrum® Control does not collect information about the size of GPFS snapshots in IBM Spectrum Scale. To view that information, you must enable IBM Spectrum Control to collect information about the size of snapshots.

About this task

Restriction: When you enable data to be collected about the size of GPFS snapshots, the collection might put a significant load on the IBM Spectrum Scale cluster during the probe. This load might affect other applications that use the same cluster.

Procedure

1. Run the following command to set the `Probe.GetGPFSnapshotSize` property for the Device server to true:

```
tpctool setdscfg -user user_ID -pwd password -url localhost:9550  
-property Probe.GetGPFSnapshotSize true
```

Where `user_ID` is an IBM Spectrum Control user ID and `password` is the password for the IBM Spectrum Control user ID.

2. Restart the IBM Spectrum Control Device server and the web server.

Alerting

Specify conditions that trigger alerts and the actions to take when those alerts are triggered, such as notify an email address. Use alert policies to define those alert conditions and notification settings for a group of resources.

Alerting functions examine the attributes, capacity, and performance of resources. If the conditions that are defined for alerts are met, the actions that are specified for the alert are taken. Typically, the actions include sending a notification. For example, if the status of a SAN Volume Controller storage system changes to Error, an alert is displayed in the Alerts page in the GUI, and an email might be sent to a storage administrator.

You can manage alerts in your storage environment in the following ways:

- Use alert policies to manage the alert definitions and notification settings that apply to different sets of resources. For example, you can use one alert policy for the servers in your test environment, and another for the servers in your production environment. Here are some important points about alert policies:
 - Alert policies manage one type of resource only. For example, if you have SAN Volume Controller and FlashSystem 900 storage systems in your storage environment, you cannot have both types of resource in one alert policy.
 - A resource can be managed by only one alert policy.
 - When you add a resource to be monitored by IBM Spectrum® Control, it is added to a default alert policy automatically. However, default alert policies are not provided for agentless servers.
 - If a resource is managed by a policy, the resource cannot have alert definitions and notification settings that are independent of the policy. The alert definitions and notification settings that apply to the resource come from the policy.

- It is not a requirement for resources to be managed by an alert policy. A resource can have its own alert definitions and notification settings, independent of an alert policy.
 - When you add a resource to be monitored by IBM Spectrum Control, it is added to a default alert policy automatically.
 - Default policies with alerts already configured are available. You can create copies of the default policies and assign resources to the new policies. Your alerts are configured with the default settings.
 - If you add a resource to a policy, any existing alert definitions for the resource are replaced by the alert definitions in that policy. You can't restore the original alert definitions for a resource after you move it into the policy.
To store a copy of the original alert definitions for a resource before you move it into a policy, create a policy based on the alert configuration of the resource. Then, move the resource to another policy. If you want to reapply the original alert definitions later, you can add the resource to the policy with the original alert configuration.
- Define alert conditions and notification settings for individual resources. It is not a requirement for resources to be managed by an alert policy. A resource can have its own alert definitions and notification settings, independent of an alert policy.
 - Define alerts and notification settings for applications and general groups. Use applications or general groups to manage alerts for groups of resource components such as volumes or pools. For example, you might want to define alerts on the response time for volumes in an application, depending on the response time requirements of the application. In this case, it is not useful to configure volume response time thresholds for the entire storage system because the storage system might serve many different applications with different needs.
 - [How alerts work](#)
Alerting functions examine the attributes, capacity, and performance of resources. If the conditions that are defined for alerts are met, the actions that are specified for the alert are taken. Typically, the actions include sending a notification. For example, if the status of a SAN Volume Controller storage system changes to Error, an alert is displayed in the Alerts page in the GUI, and an email might be sent to a storage administrator.
 - [Viewing and administering alerts](#)
View and administer the alert that were detected on monitored resources, applications, and general groups.
 - [Viewing and administering alert definitions](#)
View and administer the alert definitions and notification settings for alert policies, resources, applications, and general groups.
 - [Alert policies](#)
Use the Alert Policies page as a central location to view and manage all your alert policies.
 - [Defining notification settings for alerts](#)
You can define the alert notification settings to determine the actions that are taken when alert conditions are detected for a resource. The settings are applied to all of the alert definitions that are specified for the resource.
 - [Alert severities](#)
IBM Spectrum Control determines the severity of alert conditions that it detects on monitored resources. Use the severity level to help determine the priority in which you resolve alerts.
 - [Defining alerts](#)
Define how and when you are alerted to changes in your storage environment. You can define alerts for alert policies, for individual resources, or for sets of resources that are included in an application or general group.
 - [Configuring alert notifications](#)
Alerts can define notification actions that send email, generate Simple Network Management Protocol (SNMP) traps, or generate IBM® Tivoli® Netcool®/OMNIBus events. To enable these notification actions, you must configure IBM Spectrum Control for email, SNMP, or Tivoli Netcool/OMNIBus alert notifications.
 - [Triggering conditions for alerts](#)
Define alerts so that IBM Spectrum Control automatically notifies you when certain conditions or events are detected on monitored resources. Such conditions are the *triggering conditions* for the alert. The specific conditions that can trigger alerts depend on the type of resource that is being monitored.
 - [Alert notifications and actions](#)
Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Related tasks

- [Adding resources](#)
- [Viewing capacity alerts and violations](#)
- [Defining alert definitions for general attributes and capacity changes](#)

Related reference

- [Triggering conditions for alerts](#)

How alerts work

Alerting functions examine the attributes, capacity, and performance of resources. If the conditions that are defined for alerts are met, the actions that are specified for the alert are taken. Typically, the actions include sending a notification. For example, if the status of a SAN Volume Controller storage system changes to Error, an alert is displayed in the Alerts page in the GUI, and an email might be sent to a storage administrator.

Triggering conditions for alerts

[Learn more](#) about the conditions that can trigger alerts for each type of resource.


Event processing

Conditions that generate alerts are detected during data collection and event processing. By default, probes collect storage system data once every 24 hours. For some storage systems such as IBM Spectrum Accelerate and the XIV®, events are polled every minute from the resource. For IBM Spectrum Scale, status change events are polled frequently, typically within minutes. For other resources, events are subscription-based, where the resource itself or a data source such as a CIM agent sends the events to IBM Spectrum Control when conditions change on the resource.

Examples of storage systems that use subscription-based event processing include SAN Volume Controller and IBM FlashSystem® devices that run IBM Spectrum Virtualize. For these storage systems, a probe is automatically run when many events are received from the storage system in a short time period. To avoid performance bottlenecks, probes are run only every 20 minutes.

Determining which type of alert to use

To determine whether to define alerts in alert policies, for individual resources, or for the set of resources that are included in an application or general group, follow these guidelines:

Which type of alerts to use?	Scenario	 Learn more
Alerts defined in alert policies	You want to manage alert conditions and notification settings for a group of resources of the same type. For example, if you have several SAN Volume Controller storage systems in your environment, you can create an alert policy so that the alert definitions are the same for all of the SAN Volume Controller systems. If you have some SAN Volume Controller systems in a test environment, and some in a production environment, you can use one alert policy for the test environment, and another for the production environment.	<ul style="list-style-type: none"> • Alerts and alert policies
Resource alerts	You want to receive alert notifications about changes for a specific resource, or its internal resources. For example, for a storage system, you can alert on the attributes of the system itself, and on the attributes of its volumes, pools, ports, and other internal resources. If you define an alert for a resource, for example, a performance alert for the ports on a storage system, the alert threshold value applies to all of the ports on the storage system. You cannot apply different alert thresholds to internal resources of the same type on a resource.	<ul style="list-style-type: none"> • Defining alert definitions for performance changes • Defining custom alerts for resources
Application alerts	Use application alerts in the following scenarios: <ul style="list-style-type: none"> • You want to receive alert notifications for all the resources of a certain type in an application. For example, if your application uses multiple storage systems, you can define the storage system alerts once for the application and the alerts apply to all the storage systems. If you later add more storage systems to the application, the existing application alerts apply to those storage systems also. • You want to apply different thresholds to internal resources of the same type on a storage system. For example, you have production applications and test applications that use volumes on a SAN Volume Controller. The production applications require response times of 6 milliseconds or less while the test applications can tolerate response times up to 30 milliseconds. You can use application alerts to set separate response time thresholds for volumes used by the different applications, depending on the needs of that application. 	<ul style="list-style-type: none"> • Application alerts for performance metrics • Defining custom alerts for applications
General group alerts	Use general group alerts in the following scenarios: <ul style="list-style-type: none"> • You want to receive alert notifications about changes for a subset of the resources of a particular type. For example, you can detect when the ports that are used for replication on your SAN Volume Controller have insufficient buffer-to-buffer credit. Alert notifications are not generated for ports that are not used for replication. • You want to receive alert notifications about changes for a group of resources that are logically related. You can group all the storage systems at a specific location or all the servers that use a particular operating system. For example, you can receive alert notifications when the used capacity of any of your Linux® servers exceeds 80%. 	<ul style="list-style-type: none"> • General group alerts for performance metrics • Defining custom alerts for general groups


Tip: If a resource is in both an alert policy and a general group, the alert definitions for both the policy and the group are applied.

Viewing and administering alerts

View and administer the alert that were detected on monitored resources, applications, and general groups.

Table 1. Viewing and administering alerts

Actions	Steps
Remove alerts.	<ol style="list-style-type: none"> 1. For resources, go to the resource list page for the resource. For applications, go to Groups > Applications. For general groups, go to Groups > General Groups. 2. Right-click a resource, application, or general group and click View Details. 3. Click Alerts in the General section. 4. Right-click one or more alerts and click Remove alerts. 5. Click Remove. <p>To remove all alerts go to Home > Alerts, and then click Remove all alerts in the Actions menu.</p>
Acknowledge alerts.	<ol style="list-style-type: none"> 1. For resources, go to the resource list page for the resource. For applications, go to Groups > Applications. For general groups, go to Groups > General Groups. 2. Right-click a resource, application, or general group and click View Details. 3. Click Alerts in the General section. 4. Double-click the alert then click Acknowledge in the alert details pane. 5. To acknowledge multiple alerts, press Ctrl and click the alerts. Right-click the alerts and click View Alerts. Review the summary details of the alerts in the details pane, then click Acknowledge. <p>To acknowledge all alerts, go to Home > Alerts, and then click Acknowledge all alerts in the Actions menu.</p>

Actions	Steps
View all the alerts for resources, applications, and groups.	Go to Home > Alerts.
View all the alerts for a specific resource type.	<ol style="list-style-type: none"> 1. Go to the resource list page for the resource type that you want to view. For example, go to Network > Switches. 2. Click the Alerts tab.
View the alerts for a specific resource.	<ol style="list-style-type: none"> 1. Go to the resource list page for the resource that you want to view. For example, go to Storage > Block Storage Systems. 2. Right-click a resource and click View Details. 3. Click Alerts in the General section.
View the alerts for a specific application.	<ol style="list-style-type: none"> 1. Go to Groups > Applications. 2. Right-click an application and click View Details. 3. Click Alerts in the General section.
View the alerts for a specific general group.	<ol style="list-style-type: none"> 1. Go to Groups > General Groups. 2. Right-click a general group and click View Details. 3. Click Alerts in the General section.
View the details of an alert.	<p>Double-click the alert. The alert details are displayed in a pane.</p> <p>Click another alert to display its details. Alternatively, click the close icon .</p>
Create an alert policy. Learn more	<ol style="list-style-type: none"> 1. Go to Settings > Alert Policies . 2. To create a policy with default alert definitions, click Create Policy. 3. To create a policy by copying an existing policy, select the policy to copy, then click Actions > Copy Policy.
Modify an alert policy. Learn more	<ol style="list-style-type: none"> 1. Go to Settings > Alert Policies . 2. To modify a policy, double-click the policy.

Viewing and administering alert definitions

View and administer the alert definitions and notification settings for alert policies, resources, applications, and general groups.

Viewing and administering alert definitions

Table 1. Viewing and administering alert definitions

Actions	Steps
Create alert definitions in an alert policy. Learn more	<ol style="list-style-type: none"> 1. Go to Settings > Alert Policies. 2. To create a policy with default alert definitions, click Create Policy. 3. To create a policy by copying an existing policy, select the policy to copy, then click Actions > Copy Policy.
Modify alert definitions and notification settings in an alert policy. Learn more	<ol style="list-style-type: none"> 1. Go to Settings > Alert Policies. 2. To modify a policy, double-click the policy. 3. To edit the alert definitions, click Edit Alert Definitions on the Alert Definitions tab. 4. To edit the notification settings, click Edit Policy Notifications. <p>Restriction: You cannot modify the default alert policies.</p>
Define and edit the alert definitions or notification settings for a resource that is not managed by an alert policy. Learn more about defining alerts for resources Learn more about modifying notification settings	<ol style="list-style-type: none"> 1. Go to the list page for the resource. 2. Double-click the resource for which you want to define alerts. 3. To edit the alert definitions, click Alert Definitions in the General section of the resource details page, then click Edit Alert Definitions. 4. To edit the notification settings, click Edit Notifications.
Define and edit the alert definitions for an application or general group. Learn more	<ol style="list-style-type: none"> 1. For applications, go to Groups > Applications. For general groups, go to Groups > General Groups. 2. Right-click an application or general group and click View Alert Definitions.
Disable an alert definition for a resource, application, or general group.	<ol style="list-style-type: none"> 1. Go to the list page for the resource. 2. Right-click the resource for which you want to disable an alert, then click View Alert Definitions. 3. Click Edit Alert Definitions. 4. Click the check mark to remove the check mark from the alert definition. 5. Click Save Changes.

Alert policies

Use the Alert Policies page as a central location to view and manage all your alert policies.

To view the Alerts Policies page, click Settings > Alert Policies.

The Alert Policies page lists the default policies for your monitored resources and custom policies that you create. The page shows the following information for each alert policy:

- The types of resource that are allowed in the policy.
- The number of resources that are managed. To see which policies you are using, you can sort the list of policies by the number of resources.
- The number of alerts that are defined.
- Any email addresses that are specified to be notified about alerts instead of the default policy contacts.

You can manage the alert definitions and notification settings of the following types of resource using policies:

- Block storage systems
- File storage systems
- Object storage systems
- Hypervisors
- Switches
- Fabrics
- Servers

Here are some important points about alert policies:

- Alert policies manage one type of resource only.
- A resource can be managed by only one alert policy.
- It is not a requirement for resources to be managed by an alert policy. A resource can have its own alert definitions and notification settings, independent of an alert policy.
- If a resource is managed by a policy, the resource cannot have alert definitions and notification settings that are independent of the policy. The alert definitions and notification settings that apply to the resource come from the policy.
- When you add a resource to be monitored by IBM Spectrum® Control, it is added to a default alert policy automatically.
- Default policies with alerts already configured are available. You can create copies of the default policies and assign resources to the new policies. Your alerts are configured with the default settings.
- If you add a resource to a policy, any existing alert definitions for the resource are replaced by the alert definitions in that policy. You can't restore the original alert definitions for a resource after you move it into the policy.
To store a copy of the original alert definitions for a resource before you move it into a policy, create a policy based on the alert configuration of the resource. Then, move the resource to another policy. If you want to reapply the original alert definitions later, you can add the resource to the policy with the original alert configuration.

The Resources by Policy page lists the resources and the policies that manage the resources.

Tips:

- If you have both Cisco and Brocade switches in your storage environment, you can manage the alert definitions for the different types of switch with different policies. Create the policies from the existing alert definitions in your switches. Alternatively, create copies of the default switch policy and edit the definitions to suit your requirements.
- Default alert policies are not provided for agentless servers.

[Video] How to create alert policies

Watch a short video about how create an alert policy and view alert notifications in IBM® Storage Insights Pro.



Defining notification settings for alerts

You can define the alert notification settings to determine the actions that are taken when alert conditions are detected for a resource. The settings are applied to all of the alert definitions that are specified for the resource.

Procedure

1. To define notification settings for alerts, choose one of the following options:

Option	Steps
--------	-------

Option	Steps
Define notification settings for alerts in a policy	a. Go to Settings > Alert Policies. b. Double-click the policy. c. Click Edit Policy Notifications.
Define notification settings for alerts for a resource that is not managed by a policy	a. Go to the resource list page for the resource. For example, to modify notification settings for alerts for a block storage system, go to Storage > Block Storage Systems. To modify notification settings for a switch, go to Network > Switches. b. Right-click the resource for which you want to define alerts, then click View Alert Definitions. c. Click Edit Notifications in the Resource Notifications area.

2. If you want to send email notifications of alert violations, enter the email addresses in the Email Addresses field.

Tip: If you enter an email address in the Email Addresses field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to Settings > Notification Settings.

3. To specify other notification actions, click the switch icon  for the type of action. Enter any details required for the action.

You can specify that the following notification actions are taken when alert conditions are detected on monitored resources:

Run script

Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks.

[Learn more.](#)

Restriction: You can only specify to run a script for an individual alert; you cannot specify to run a script for all alerts in an alert policy.

SNMP

Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.

Netcool® / OMNIBus

Send alert notifications to a Netcool server or OMNIBus EIF probe server within your environment that was configured to receive IBM Spectrum Control alerts.

Windows event log or UNIX syslog

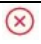





Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

4. Click Save Changes.

Alert severities

IBM Spectrum® Control determines the severity of alert conditions that it detects on monitored resources. Use the severity level to help determine the priority in which you resolve alerts.

Each severity is represented by an icon in the web-based GUI.

Icon	Status	Description
	Critical	An alert with a Critical severity represents a serious problem on a resource or on its internal resources. Resolve these problems as soon as possible. Review the condition that triggered an alert for more information about the problem.
	Critical - Acknowledged	An alert with a Critical severity was acknowledged. A Critical - Acknowledged severity indicates that an alert was reviewed and is either resolved or can be ignored.
	Warning	An alert with a Warning severity represents potential problems on a resource or on its internal resources. Resolve these problems after you fix any critical alerts. Review the condition that triggered an alert for more information about the problem.
	Warning - Acknowledged	An alert with a Warning severity was acknowledged. A Warning - Acknowledged severity indicates that an alert was reviewed and is either resolved or can be ignored.
	Informational	An alert with an Informational severity does not represent a problem, but is intended to provide information about actions related to a resource.
	Informational - Acknowledged	An alert with an Informational severity was acknowledged. An Informational - Acknowledged severity indicates that an alert was reviewed and can be ignored.

Defining alerts

Define how and when you are alerted to changes in your storage environment. You can define alerts for alert policies, for individual resources, or for sets of resources that are included in an application or general group.

Before you begin

Configure IBM Spectrum® Control to send alert notifications by using email, SNMP traps, or a Tivoli® Netcool®/OMNIBus server.

- [Defining alerts for resources](#)**
Define alerts for changes in the configuration, attributes, and performance of monitored resources.
- [Defining alerts for applications](#)**
Define alerts for changes in the configuration, attributes, and performance of the servers , volumes, filesets, and shares in your application.
- [Defining alerts for general groups](#)**
Define alerts for changes in the configuration, attributes, and performance of the resources in your general groups.

Related reference

- [Alert notifications and actions](#)

Defining alerts for resources

Define alerts for changes in the configuration, attributes, and performance of monitored resources.

- [Defining alert definitions for general attributes and capacity changes](#)
You can define alerts that are triggered when the attributes or capacity of a resource changes. Attributes represent the key properties and configuration of a resource, such as status, versions, removals, discoveries, and data collection status. Capacity represents storage statistics such as available capacity, used capacity, drive capacity, reserved capacity, and more.
- [Defining alert definitions for performance changes](#)
You can define alerts that are triggered when the performance of a resource falls outside a specified threshold.
- [Defining custom alerts for resources](#)
You can define alerts that are triggered when two or more changes occur in the attributes, capacity, and performance of resources.

Defining alert definitions for general attributes and capacity changes

You can define alerts that are triggered when the attributes or capacity of a resource changes. Attributes represent the key properties and configuration of a resource, such as status, versions, removals, discoveries, and data collection status. Capacity represents storage statistics such as available capacity, used capacity, drive capacity, reserved capacity, and more.

About this task

Asset, capacity, and configuration metadata are aggregated and collected when probes collect the metadata from the resources. By default, the metadata that is collected from storage systems is refreshed every 24 hours. Define alerts to track daily changes to the attributes of a storage system.

Procedure

1. To define alerts for resources, choose one of the following options:

Option	Steps
Define alerts for a policy	<ol style="list-style-type: none"> a. Go to Settings > Alert Policies. b. Double-click the policy. c. Click Edit Alert Definitions on the Alert Definitions tab.
Define alerts for a resource that is not managed by a policy	<ol style="list-style-type: none"> a. Go to the resource list page for the resource. For example, to define alerts for a block storage system, go to Storage > Block Storage Systems. To define alerts for a switch, go to Network > Switches. b. Right-click the resource for which you want to define alerts, then click View Alert Definitions. c. Click Edit Alert Definitions.

2. Click the type of resource that you want to alert on. For example, click Storage System.

3. Click the category of the attributes that you want to alert on.

Category	Description
General	Attributes for the key properties of a resource, such as status, version changes, removals, discoveries, state, and data collection status.
Capacity	Attributes for capacity statistics of a resource, such as available capacity, used capacity, drive capacity, reserved capacity, and more.

Restriction: Not all categories are available for some resources.

4. To enable the alert for an attribute, click the check mark for the attribute.
The advanced options for the alert, such as notification frequency, are displayed.
5. Specify the conditions for generating an alert for an attribute.
Conditions can include operators such as >=, or <=. Conditions can also include storage values and time values.
For example, for a capacity attribute such as Available Capacity, you can specify that an alert is generated when the amount of available capacity on a resource's pools is less than or equal to 50 GiB.




☒ Available Capacity

Operator	Value	Unit
<=	50	GiB

Tips:

- Not all attributes require conditions to generate an alert. For example, you can enable an alert for the Deleted Volume attribute, but you don't need to specify any conditions.
 - Some attributes use operators such as *is*, *is not*, *contains*, and *changes*. For example, for the Firmware attribute for a DS8000® you can select the operator Contains and enter R5 in the value field. An alert is triggered if the firmware is at the R5 level rather than at a later version such as R6.1, R6.2, or R6.3. You can use this alert definition if you want to be notified when the firmware for a storage system is reverted to a previous version.
6. Assign a severity to an alert.

Assigning a severity can help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
 Critical	Alert is critical and needs to be resolved. For example, alerts that notify you when the amount of available capacity on a file system falls below a specified threshold.
 Warning	Alerts that are not critical, but represent potential problems. For example, alerts that notify you when the status of a data collection job is not normal.
 Informational	Alerts that might not require any action to resolve and are primarily for informational purposes. For example, alerts that are generated when a new pool is added to a storage system.

7. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the Email Override field.
Tip: If you enter an email address in the Email Override field, only that email address receives notifications for the alert. The following contacts do not receive notifications:
 - Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
 - Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to [Settings > Notification Settings](#).
8. Optional: Click View Additional Options to specify how frequently you are notified of alerts.
Use these settings to avoid triggering too many alerts for some conditions.
9. Optional: Click View Additional Options to specify that the following actions are taken when alert conditions are detected on monitored resources:

Run script

Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks.

[Learn more](#)

Netcool® / OMNIBus

Send alert notifications to a Netcool server or OMNIBus EIF probe server within your environment that was configured to receive IBM Spectrum Control alerts.

SNMP

Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.



Windows event log or UNIX syslog

Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

10. Optional: [+](#) Duplicate an alert.

Use this action when you want to define another alert for the same attribute but with different conditions and settings.

Duplicating alerts can be helpful in the following situations:

- When you want to generate separate warning alerts and critical alerts for different conditions on the same attribute.
For example, for a capacity attribute such as Available Capacity, you might want to define the following alerts:
 - Define a warning alert  to be generated when the amount of available capacity on a resource's pools is less than or equal to 50 GiB.
 - Duplicate the alert, but this time, specify a critical severity  when the amount of available capacity on a resource's pools is less than or equal to 10 GiB.
- When you want to send, alert notifications to different people based on the severity of an alert.
In the previous example for the Available Pool Space attribute, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators.

11. Click Save Changes.

Results

To view all the alerts generated by IBM Spectrum Control, go to [Home > Alerts](#) in the GUI.

Related concepts

- [Collecting data](#)
- [Alerting](#)

Related tasks

- [Viewing capacity alerts and violations](#)

Related reference

- [Alert notifications and actions](#)

Defining alert definitions for performance changes

You can define alerts that are triggered when the performance of a resource falls outside a specified threshold.

Procedure

1. To define alerts for resources, choose one of the following options:

Option	Steps
Define alerts for a policy	a. Go to Settings > Alert Policies . b. Double-click the policy. c. Click Edit Alert Definitions on the Alert Definitions tab.
Define alerts for a resource that is not managed by a policy	a. Go to the resource list page for the resource. For example, to define alerts for a block storage system, go to Storage > Block Storage Systems . To define alerts for a switch, go to Network > Switches . b. Right-click the resource for which you want to define alerts, then click View Alert Definitions. c. Click Edit Alert Definitions.

2. Click the type of resource that you want to alert on. For example, click Storage System.
3. Click the Performance category.




4. To enable the alert for a performance metric, click the check mark for the metric. If the metric that you want is not displayed, click Add Metrics then select the metric you want.
5. Specify the conditions for generating an alert.
Conditions include an operator and a threshold value.
 - a. Select an operator.
An operator determines whether an alert is triggered when the performance of a resource is *greater than or equal to* or *less than or equal to* the specified threshold value.
 - b. Enter a threshold value.
For example, to trigger an alert if the Total I/O Rate for a storage system is greater than or equal to 500 ops/s, enter the value 500.
Tips for threshold values:
 - IBM Spectrum® Control provides recommended values for threshold values that do not vary much between environments. For example, the default threshold values for Port Send Bandwidth Percentage are greater than or equal to 75% for warning alerts, and greater than or equal to 85% for critical alerts.
However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
 - For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
6. Optional: Click View Performance to view a chart of the performance of the resource. Use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.
The chart displays a horizontal color line at the specified threshold value. The color of the line indicates the severity of the alert:
 - Critical alert: red
 - Warning alert: yellow
 - Information alert: blue

For multi-conditional alerts, the chart displays a horizontal line for each condition that shows the threshold value and severity.

To customize the chart, click Top 10 or Bottom 10 to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.

7. Assign a severity to an alert.

Assigning a severity can help you more quickly to identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
 Critical	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%.
 Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
 Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

8. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the Email Override field.
Tip: If you enter an email address in the Email Override field, only that email address receives notifications for the alert. The following contacts do not receive notifications:
 - Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
 - Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to Settings > Notification Settings.
9. Optional: Click View Additional Options to specify how frequently you are notified of alerts.
Use these settings to avoid triggering too many alerts for some conditions.
10. Optional: Click View Additional Options to specify that the following actions are taken when alert conditions are detected on monitored resources:

Run script

Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks.
[Learn more.](#)

Netcool® / OMNIBus

Send alert notifications to a Netcool server or OMNIBus EIF probe server within your environment that was configured to receive IBM Spectrum Control alerts.

SNMP

Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.



Windows event log or UNIX syslog

Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

11. Optional:  Duplicate an alert.

Use this action when you want to define another alert for the same metric but with different conditions and settings.

Duplicating alerts can be helpful in the following situations:

- When you want to generate separate warning alerts and critical alerts for different thresholds on the same metric.
For example, for the CRC Error Rate metric for ports, you might want to define the following alerts:
 - Define a warning alert  to be generated when the number of frames per second that are received with cyclic redundancy check (CRC) errors is greater than or equal to 0.01 counts per second.
 - Duplicate the alert, but this time, specify a critical severity  when the CRC error rate is greater than or equal to 0.03 counts per second.
- When you want to send alert notifications to different people based on the severity of an alert.
In the previous example for the CRC Error Rate metric, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

12. Click Save Changes.

Results

To view all the alerts generated by IBM Spectrum Control, go to Home > Alerts in the GUI.

Tip: If a performance monitor is already collecting data about a resource when you add, modify, or remove a performance alert for that resource, changes are applied dynamically. You do not have to stop and restart the performance monitor to apply the changes. A confirmation message is recorded in the log of the performance monitor when the alert is updated.

Related tasks

- [Collecting performance data by using IBM Spectrum Control performance monitors](#)

Related reference

- [Performance metrics](#)
- [Alert notifications and actions](#)

Defining custom alerts for resources

You can define alerts that are triggered when two or more changes occur in the attributes, capacity, and performance of resources.

About this task

To define a custom alert, select the general attributes, capacity, and performance metrics that you want to combine to trigger an alert and specify their conditions and threshold values. You can combine conditions for the resource and its internal resources into a custom alert. The alert is triggered when the conditions for the attributes and capacity of the resource are met, and the performance of the resource falls outside the threshold values.


For example, you can create a custom alert that notifies you when the overall response time for the volumes on a SAN Volume Controller system is worse than 20 milliseconds per operation *and* the system CPU utilization on the nodes on the system is greater than 70%. The Overall Response Time is a metric that measures the average number of milliseconds that it takes to service each I/O operation on a volume. The System CPU Utilization is a metric that measures the average percentage of time that the processors on nodes are busy doing system I/O tasks.

Procedure

1. To define alerts for resources, choose one of the following options:




Option	Steps
Define alerts for a policy	a. Go to Settings > Alert Policies. b. Double-click the policy. c. Click Edit Alert Definitions on the Alert Definitions tab.
Define alerts for a resource that is not managed by a policy	a. Go to the resource list page for the resource. For example, to define alerts for a block storage system, go to Storage > Block Storage Systems. To define alerts for a switch, go to Network > Switches. b. Right-click the resource for which you want to define alerts, then click View Alert Definitions. c. Click Edit Alert Definitions.

2. Click Custom.

3.  Click the create alert icon, then enter a name for the alert.

4. Assign a severity to the alert.

Assigning a severity can help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization.

Option	Description
 Critical	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%. The default severity for custom alerts is critical.
 Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
 Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

5. Select a component, category, and group for the alert.

For example, select Storage System, Capacity, and Available Capacity.

6. To generate an alert for a general or capacity attribute, specify the conditions for the alert.

Conditions can include operators such as >=, or <=. Conditions can also include storage values and time values.

For example, for a capacity attribute such as Available Capacity, you can specify that an alert is generated when the amount of available capacity on a resource's pools is less than or equal to 50 GiB.

☒ Available Capacity

Operator	Value	Unit
<div><=</div> <div>▼</div>	<div>50</div> <div>▲</div> <div>▼</div>	<div>GiB</div> <div>▼</div>


Tips:

- Not all attributes require conditions to generate an alert. For example, you can enable an alert for the Deleted Volume attribute, but you don't need to specify any conditions.
- Some attributes use operators such as *is*, *is not*, *contains*, and *changes*. For example, for the Firmware attribute for a DS8000® you can select the operator Contains and enter R5 in the value field. An alert is triggered if the firmware is at the R5 level rather than at a later version such as R6.1, R6.2, or R6.3. You can use this alert definition if you want to be notified when the firmware for a storage system is reverted to a previous version.

7. To generate an alert for a performance metric, specify the conditions for the alert.
Conditions include an operator and a threshold value.
 - a. Select an operator.
An operator determines if an alert is triggered when the performance of a resource is *greater than or equal to* or *less than or equal to* the specified threshold value.
 - b. Enter a threshold value.
For example, to trigger an alert if the Total I/O Rate for a storage system is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

 - IBM Spectrum® Control provides recommended values for threshold values that do not vary much between environments. For example, the default threshold values for Port Send Bandwidth Percentage are greater than or equal to 75% for warning alerts, and greater than or equal to 85% for critical alerts.
However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
 - For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
8. Optional: Click View Performance to view a chart of the performance of the resource. Use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.
The chart uses colored lines to represent the different threshold values and severities that can be defined for an alert:
 - Critical alert: red
 - Warning alert: orange
 - Information alert: blue

To customize the chart, click Top 10 or Bottom 10 to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.
9.  Duplicate an alert.
Use this action to add a second and subsequent condition to the alert.
10. Repeat steps 5 - 9 to add more conditions to the alert.
11. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the Email Override field.
Tip: If you enter an email address in the Email Override field, only that email address receives notifications for the alert. The following contacts do not receive notifications:
 - Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
 - Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to Settings > Notification Settings.
12. Optional: Click View Additional Options to specify how frequently you are notified of alerts.
Use these settings to avoid triggering too many alerts for some conditions.
13. Optional: Click View Additional Options to specify that the following actions are taken when alert conditions are detected on monitored resources:

Run script
Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks.
[Learn more.](#)

Netcool® / OMNIBus
Send alert notifications to a Netcool server or OMNIBus EIF probe server within your environment that was configured to receive IBM Spectrum Control alerts.

SNMP
Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.

Windows event log or UNIX syslog
Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.
14. Click Save Changes.

Results

To view all the alerts generated by IBM Spectrum Control, go to Home > Alerts in the GUI.

- [Scenarios for custom alerts](#)
Create custom alerts with one or more conditions that trigger when all the conditions are met for a given resource. By creating a custom alert, you can detect multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred in your storage or SAN fabric.

Related reference

- [Alert notifications and actions](#)

Scenarios for custom alerts

Create custom alerts with one or more conditions that trigger when all the conditions are met for a given resource. By creating a custom alert, you can detect multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred in your storage or SAN fabric.

The requirements of your environment determine the custom alerts that you create. For example, your storage systems might run critical production applications on tier 1 storage. In this case, you don't want the performance of the tier 1 storage to fall below a certain threshold. To be notified when that situation occurs, you can create a custom alert that checks if the overall response time of tier 1 storage is too high.

Tip: You can define alerts for storage systems, servers, hypervisors, fabrics, and switches and their internal resources. However, the resources in an alert must be along the same data path. For example, if you create an alert for a storage system volume, resources along the same data path include the pool that volume is part of, the switch

ports that are used to access that volume, and the host that maps the volume.
Use the following example scenarios as guides to help you create custom alerts for your environment.

- [Receive alerts when the response time of storage on a specific tier is too high](#)
- [Receive alerts when the response time of volumes is too high during times of active I/O](#)
- [Receive alerts when the response time of volumes is too high, but do not generate these alerts when batch and backup jobs are running](#)
- [Receive alerts if a port is being used for both inter-node communication and host I/O exchanges](#)
- [Receive alerts for link resets that are not associated with link initialization](#)
- [Receive alerts for invalid word transmissions that are not associated with link initialization](#)

Receive alerts when the response time of storage on a specific tier is too high

Your storage systems run both critical production applications and noncritical test applications. The production applications use tier 1 storage, while the test applications use storage on tiers 2 and 3.

To ensure consistent, top performance for tier 1 storage, you want to be notified when its response time is higher than 6 ms/op so that you can resolve the bottleneck. However, to avoid too many alerts, you do not want to receive notifications when the response time of tier 2 or 3 storage exceeds 6 ms/op.




Solution

Define a custom alert that checks if the volumes used by an application are in Tier 1 pools and if their Overall Response Time is higher than 6 ms/op.

For a storage system, set up a custom alert with the following attributes and conditions:

☒ High response time for tier 1 stor

Severity



ComponentCategoryAttribute

PoolsGeneralTier

ConditionValue

Is1

ComponentCategoryAttribute

VolumesPerformanceOverall Response Time

OperatorValue

>=6ms/op

Attribute	Condition
Overall Response time (Storage System > Volumes > Performance)	Greater than or equal to (>=) 6 ms/op
Tier (Storage System > Pools > General)	Tier 1

Receive alerts when the response time of volumes is too high during times of active I/O




You care about high read response times on your volumes, but they can be caused by cache misses when there is only a trickle of I/O.

Solution

Define a custom alert with volume-level thresholds that combines checks for response times and I/O.

☒ Response time with workload

Severity



ComponentCategoryAttribute

VolumesPerformanceOverall Response Time

OperatorValue

>=10ms/op

ComponentCategoryAttribute

VolumesPerformanceTotal I/O Rate - overall

OperatorValue

>=50ops/s

Attribute	Condition
Overall Response time (Storage System > Volumes > Performance)	Greater than or equal to (\geq) 10 ms/op
Total I/O Rate - overall (Storage System > Volumes > Performance)	Greater than or equal to (\geq) 50 ms/op

Receive alerts when the response time of volumes is too high, but do not generate these alerts when batch and backup jobs are running

To ensure the consistent, top performance of your volumes, you want to be notified when their response times are becoming too high. However, to avoid too many alerts, you do not want to receive notifications when batch and backup jobs are running on your storage. You understand that these jobs can cause an expected spike in response times and do not require action on your part.

Solution

Define a custom alert that checks if the Read Response Time of volumes exceeds an amount that is more than expected in your environment and the Read Transfer Size is less than 256 KiB/op. Typically, read transfer sizes greater than 256KiB/op indicate that batch or backup jobs are running in the background.

For a storage system, set up a custom alert with the following attributes and conditions:

☒ Slow performance of volumes

Severity

Component

Volumes

Category

Performance

Attribute

Read Response Time

Operator

\geq

Value

20

ms/op

Component

Volumes

Category

Performance

Attribute

Read Transfer Size

Operator

\leq

Value

256

KiB/op

Attribute	Condition
Read Response Time (Custom > Volumes > Performance)	Greater than or equal to (\geq) 20 ms/op
Read Transfer Size (Custom > Volumes > Performance)	Less than or equal to (\leq) 256 KiB/op

Receive alerts if a port is being used for both inter-node communication and host I/O exchanges

You want to avoid potential bottlenecks by ensuring that storage system ports aren't being used for both inter-node communication in the local cluster and for I/O exchanges to host computers. You can also use this custom alert to check for adherence to best practices that are related to configuring ports for nodes with 8 or more ports. It does not apply to nodes that contain only 4 ports.

Solution

Define a custom alert that checks if the I/O rate for ports indicates exchanges between local nodes and hosts. For a storage system, set up a custom alert with the following attributes and conditions:

☒ Dual use ports

Severity

Component

FC Ports

Category

Performance

Attribute

Total Port to Host I/O...

Operator

\geq

Value

0.01

ops/s

Component

FC Ports

Category

Performance

Attribute

Total Port to Local No...

Operator

\geq

Value

0.01

ops/s

Attribute	Condition
Total Port-to-Host I/O Rate (Custom > Ports > Performance)	Greater than or equal to (\geq) .01 ops/s
Total Port-to-Local Node I/O Rate (Custom > Ports > Performance)	Greater than or equal to (\geq) .01 ops/s

Tip: Optionally, you can define other custom alerts to be notified of this situation, depending on your storage requirements. For example:

Attribute	Condition
Total Port-to-Disk I/O Rate (Custom > Ports > Performance)	Greater than or equal to (\geq) .01 ops/s
Total Port-to-Local Node I/O Rate (Custom > Ports > Performance)	Greater than or equal to (\geq) .01 ops/s
Total Port-to-Remote I/O Rate (Custom > Ports > Performance)	Greater than or equal to (\geq) .01 ops/s


Receive alerts for link resets that are not associated with link initialization

You want to identify link resets that are generated in response to hardware failures or link congestion. Link Resets generated by link initialization are ignored.

Solution

Define a custom alert that checks if link resets occur and if those resets are not associated with a link initialization. For a switch, set up a custom alert with the following attributes and conditions:

☒ Link reset not initialized

Severity


Component FC Ports **Category** Performance **Attribute** Loss of Sync Rate

Operator \leq **Value** 0 cnt/s

Component FC Ports **Category** Performance **Attribute** Loss of Signal Rate

Operator \leq **Value** 0 cnt/s

Attribute	Condition
Link Reset Received Rate (Storage System > Ports > Performance)	Greater than or equal to (\geq) .01 cnt/s
OR	
Link Transmitted Received Rate (Storage System > Ports > Performance)	
Sync Loss (Storage System > Ports > Performance)	Less than or equal to (\leq) 0 cnt/s
Signal Loss (Storage System > Ports > Performance)	Less than or equal to (\leq) 0 cnt/s

Receive alerts for invalid word transmissions that are not associated with link initialization

You want to identify invalid transmission words that are generated because of poor link quality. Poor or marginal link quality can be caused by a bad SFP, HBA, or cable. Invalid transmission words that are generated by link initialization are ignored.

Solution

Define a custom alert that checks if invalid word transmissions occur and if those resets are not associated with a link initialization. For a switch, set up a custom alert with the following attributes and conditions:

☒ Invalid word transmission not ini

Severity

Component

FC Ports

Category

Performance

Attribute

Loss of Sync Rate

Operator

<=

Value

0

cnt/s

-

+

Component

FC Ports

Category

Performance

Attribute

Loss of Signal Rate

Operator

<=

Value

0

cnt/s

-

+

Component

FC Ports

Category

Performance

Attribute

Invalid Transmission ...

Operator

>=

Value

0.01

cnt/s

-

+

Attribute	Condition
Invalid Transmission Word Rate (Storage System > Ports > Performance)	Greater than or equal to (>=) .01 cnt/s
Sync Loss (Storage System > Ports > Performance)	Less than or equal to (<=) 0 cnt/s
Signal Loss (Storage System > Ports > Performance)	Less than or equal to (<=) 0 cnt/s

Defining alerts for applications

Define alerts for changes in the configuration, attributes, and performance of the servers , volumes, filesets, and shares in your application.

About this task

You can quickly identify and address the critical conditions that are detected on the servers , volumes, filesets, and shares in your application by using alerts. Application alerts can be helpful in the following situations:

- When you want to apply different thresholds to internal resources of the same type on a storage system.
For example, you have production applications and test applications that use volumes on a SAN Volume Controller. The production applications require response times of 6 milliseconds or less while the test applications can tolerate response times up to 30 milliseconds. You can use application alerts to set separate response time thresholds for volumes used by the different applications, depending on the needs of that application.
- When you want to quickly define alerts for multiple resources of the same type. You can define alerts once for the application and the alerts apply to all the resources of that type in the application.
For example, if your application uses multiple servers , you can define the servers alerts once for the application and the alerts apply to all the servers . If you later add more servers to the application, the existing application alerts apply to those servers also.

Tips for working with alerts in application hierarchies:

- When you define an application alert for a resource such as a volume, the alert applies to all the resources of that type that belong to that application and all child applications.
- If you define an application alert for a resource such as a servers , and the servers also belongs to child applications, the alert is generated once at the parent application level. Separate alerts are not generated for each of the child applications that contain the servers .
- If a child application has the same alert as a parent application but with different conditions, separate alerts are generated for the child application for the different alert conditions.
- Defining application alerts for attribute and capacity changes**
You can define alerts that are triggered when the attributes or capacity of the servers , volumes, filesets, or shares in an application change.
- Defining application alerts for performance metrics**
You can define alerts that are triggered when the performance of the volumes that belong to an application fall outside a specified threshold.
- Defining custom alerts for applications**
Define a custom alert to combine multiple conditions from multiple resources for an application in a single alert. By creating a custom alert, you can analyze multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred on the servers, volumes, filesets, and shares in an application.

Related concepts

- [Applications](#)

Defining application alerts for attribute and capacity changes

You can define alerts that are triggered when the attributes or capacity of the servers , volumes, filesets, or shares in an application change.

About this task

Attributes represent the key properties and configuration of a resource, such as status, removals, discoveries, and data collection status. Capacity represents storage statistics such as available and used file system capacity, drive capacity, and volume used capacity.

Procedure

To define an application alert for attribute and capacity changes, complete these steps:

1. In the menu bar, select Groups, Applications.
2. Right-click an application in the list and click View Alert Definitions.
3. Click Edit Alert Definitions.
4. Click the type of resource that you want to set an alert for. You can set an alert for one or more of the attributes of servers , volumes, filesets, or shares.
5. Click the category of the attributes that you want to set an alert for.

Category	Description
General	Attributes for the key properties of a resource, such as status, removals, discoveries, and data collection status.
Capacity	Attributes for capacity statistics of a resource, such as used capacity, available file system capacity, drive capacity, and volume used capacity.




6. To enable the alert for an attribute, click the check mark for the attribute.
7. Specify the conditions for generating an alert for an attribute.
Conditions can include operators such as *greater than or equal to*, or *less than or equal to*. Conditions can also include storage values and time values.
For example, for a capacity attribute such as Available Capacity, you can be alerted when the amount of available capacity on a volume is less than or equal to 50 GiB.

☒ Available Capacity

Operator	Value	Unit
<div><=</div>	<div>50</div>	<div>GiB</div>

Tip: Not all attributes require conditions to generate an alert. For example, you can enable an alert for the New Fileset attribute for an IBM Spectrum Scale, but you don't need to specify any conditions.

8. Assign a severity to an alert to help you more quickly identify and address the critical conditions that are detected on resources.
The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
 Critical	Assign this severity to alerts that are critical and must be resolved. For example, assign a critical severity to alerts that notify you when the amount of available capacity on a file system falls below a specified threshold.
 Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the status of a data collection job is not normal.
 Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes. For example, assign an informational severity to alerts that are generated when a new fileset is added to a storage system.

9. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the Email Override field.

Tip: If you enter an email address in the Email Override field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to Settings, Notification Settings.

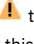

10. Optional: Click View Additional Options to specify how frequently you are notified of alerts.

Use these settings to avoid triggering too many alerts for some conditions.

11. Optional: Click the icon  to duplicate an alert.

Use this action when you want to define another alert for the same attribute but with different criteria and settings.

Duplicating alerts can be helpful in the following situations:

- When you want to generate separate warning alerts and critical alerts for different conditions on the same attribute.
For example, for a capacity attribute such as Available Capacity, you might want to define the following alerts:
 - Define a warning alert  to be generated when the amount of available capacity on a resource's disks is less than or equal to 50 GiB.
 - Duplicate the alert, but this time, specify a critical severity  when the amount of available capacity on a resource's disks is less than or equal to 10 GiB.
- When you want to send alert notifications to different people based on the severity of an alert.
In the previous example for the Available Capacity attribute, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

12. Click Save Changes.

Related reference

- [Alert notifications and actions](#)

Defining application alerts for performance metrics

You can define alerts that are triggered when the performance of the volumes that belong to an application fall outside a specified threshold.

About this task

To define an application alert for performance metrics, select the metric that you want to measure for the volumes in an application and specify a threshold value. When the performance of a volume falls outside the threshold, an alert is generated. For example, by using the Read Response Time metric you can define an alert that notifies you when the response time for a volume exceeds 20 milliseconds per read operation.

Procedure

To define an application alert for performance metrics, complete these steps:

1. In the menu bar, select Groups > Applications.
2. Right-click an application in the list and click View Alert Definitions.
3. Click Edit Alert Definitions.
4. Click Volumes and then click Performance.
5. Click Add Metrics.
6. Select one or more metrics to alert on and click OK.
7. To enable the alert for a performance metric, click the check mark for the metric.
8. Specify the conditions for generating an alert.

Conditions include an operator and a threshold value.

- a. Select an operator.

An operator determines if an alert is triggered when the performance of a resource is *greater than or equal to* or *less than or equal to* the specified threshold value.

- b. Enter a threshold value.




For example, to trigger an alert if the Total I/O Rate for a storage system is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

- IBM Spectrum® Control provides recommended values for threshold values that do not vary much between environments. For example, the default threshold values for Port Send Bandwidth Percentage are greater than or equal to 75% for warning alerts, and greater than or equal to 85% for critical alerts.
However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.

9. Assign a severity to an alert.

Assigning a severity can help you more quickly identify and address the critical performance conditions that are detected on volumes. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
 Critical	Assign this severity to alerts that are critical and must be resolved. For example, you might assign a critical severity to alerts that notify you when the average Read Response Time for volumes is greater than or equal to 20 ms/op.
 Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, you might assign a warning severity to alerts that notify you when the average Read Response Time for volumes is greater than 10 ms/op but less than 20 ms/op.
 Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

10. Optional: Click View Performance to view a chart of the performance of the resource. Use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.

The chart uses colored lines to represent the different threshold values and severities that can be defined for an alert:

- Critical alert: red
- Warning alert: orange
- Information alert: blue

To customize the chart, click Top 10 or Bottom 10 to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.


11. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the Email Override field.

Tip: If you enter an email address in the Email Override field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to Settings > Notification Settings.



12. Optional: Click View Additional Options to specify how frequently you are notified of alerts.

Use these settings to avoid triggering too many alerts for some conditions.

13. Optional: Click the icon  to duplicate an alert.

Use this action when you want to define another alert for the same metric but with different criteria and settings.

Duplicating alerts can be helpful in the following situations:

- When you want to generate separate warning alerts and critical alerts for different thresholds on the same metric.
For example, for the Read Response Time metric for volumes, you might want to define the following alerts:
 - Define a warning alert  to be generated when the average time that it takes to service each read operation is greater than or equal to 10 milliseconds.
 - Duplicate the alert, but this time, specify a critical severity  when the average time is greater than or equal to 20 milliseconds.
- When you want to send alert notifications to different people based on the severity of an alert.
In the previous example for the Read Response Time metric, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

14. Click Save Changes.

Related tasks

- [Verifying that a performance monitor is running for a resource](#)

Related reference

- [Alert notifications and actions](#)

Defining custom alerts for applications


Define a custom alert to combine multiple conditions from multiple resources for an application in a single alert. By creating a custom alert, you can analyze multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred on the servers, volumes, filesets, and shares in an application.

About this task




For example, you can create a custom alert that notifies you when the used capacity on both volumes and filesets is greater than a specified amount.

Tip: In a custom alert, the resource types that you specify conditions for must be associated with the same type of top-level resource. For example, if you include conditions for storage system volumes, you can also include additional conditions for filesets and shares, but you cannot include conditions for servers.

Procedure

1. In the menu bar, select Groups > Applications.
2. Right-click an application in the list and click View Alert Definitions.
3. Click Edit Alert Definitions.
4. Click Custom.
5. Click the  Create Alert icon, then enter a name for the alert.
6. Assign a severity to the alert.

Assigning a severity can help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization.

Option	Description
 Critical	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%. The default severity for custom alerts is critical.
 Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
 Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

7. Select a component, category, and group for the alert.
For example, select Volumes, Capacity, and Used Capacity.
8. For general and capacity attributes, specify the criteria for generating an alert.
Use criteria such as greater than or equal to, less than or equal to, storage values, and time measurements to customize the conditions under which attributes generate alerts.
For example, for a capacity attribute such as Used Capacity, you can specify that an alert is generated when the amount of used capacity on a resource is more than or equal to 75%. The operator (less than or equal to) + a specified amount of space (75) + the unit of measurement (%) is the criteria that determines if an alert is generated.
Tips:
 - Not all attributes require criteria for generating alerts. The category and type of an attribute determines whether you can specify criteria and the options that you can select.
 - Some attributes can use the operators such as *is*, *is not*, and *changes*. For example, for the Used Inodes attribute for a fileset, select the operator changes to be notified if the current number of used inodes for a fileset changes.
9. For performance attributes, click Add Metrics. Specify the criteria for generating an alert.
Criteria includes an operator and a threshold value.
 - a. Select an operator.
An operator determines if an alert is triggered when the performance of a resource is *greater than or equal to* or *less than or equal to* the specified threshold value.
 - b. Enter a threshold value.
For example, to trigger an alert if the Total I/O Rate for a volume is greater than or equal to 500 ops/s, enter the value 500.
Tips for threshold values:
 - For metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
 - For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
10. Optional: Click View Performance to view a chart of the performance of the resource. Use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.
The chart uses colored lines to represent the different threshold values and severities that can be defined for an alert:
 - Critical alert: red
 - Warning alert: orange
 - Information alert: blue
To customize the chart, click Top 10 or Bottom 10 to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.

11. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the Email Override field.
Tip: If you enter an email address in the Email Override field, only that email address receives notifications for the alert. The following contacts do not receive notifications:
 - Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
 - Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to Settings > Notification Settings.
12. Optional: Click View Additional Options to specify how frequently you are notified of alerts.
Use these settings to avoid triggering too many alerts for some conditions.
13. + Add another condition to the custom alert and repeat steps 7 - 11.
If you do not want to add another condition, continue to the next step.
14. Click Save Changes.

Related reference

- [Alert notifications and actions](#)

Defining alerts for general groups

Define alerts for changes in the configuration, attributes, and performance of the resources in your general groups.

Before you begin

Before you can define general group alerts for a resource type, you must first add resources of that type to the group or one of its subgroups. For example, if you want to set an alert for the attributes of a switch, you must first add one or more switches to your general group or one of its subgroups.

About this task

Tips for working with alerts in general group hierarchies:

- When you edit the alert definitions for a general group, you can edit the alert definitions for its subgroups at the same time. The resources for all subgroups of the general group are shown when you edit the alert definitions for a parent general group.
- If you define a general group alert for a resource such as a server, and the server also belongs to subgroups, the alert is generated once at the parent group level. Separate alerts are not generated for each of the subgroups that contain the server.
- If a subgroup has the same alert as a parent group but with different conditions, separate alerts are generated for the subgroup for the different alert conditions.
- [Defining general group alerts for attribute and capacity changes](#)
You can define alerts that are triggered when the attributes or capacity of the resources that belong to a general group change.
- [Defining general group alerts for performance metrics](#)
You can define alerts that are triggered when the performance of the storage systems or switches that belong to a general group fall outside a specified threshold.
- [Defining custom alerts for general groups](#)
Define a custom alert to combine multiple conditions from multiple resources for a general group in a single alert. By creating a custom alert, you can analyze multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred on the resources in a group.

Related reference

- [General groups](#)

Defining general group alerts for attribute and capacity changes

You can define alerts that are triggered when the attributes or capacity of the resources that belong to a general group change.

Before you begin

Attributes represent the key properties and configuration of a resource, such as status, removals, discoveries, and data collection status. Capacity represents storage statistics such as available and used file system capacity, drive capacity, and volume used capacity.

Procedure

To define a general group alert for attribute and capacity changes, complete these steps:

1. In the menu bar, select Groups > General Groups.
2. Right-click a general group in the list and click View Alert Definitions.
3. Click Edit Alert Definitions.
4. Click the type of resource that you want to set an alert for.
You can set an alert for one or more of the attributes of the resource itself or its internal resources.
For example, for a switch, you can set an alert for the attributes of the switch itself and for the attributes of the blades, inter-switch connections, and ports.
5. Click the category of the attributes that you want to alert for.

Category	Description
General	Attributes for the key properties of a resource, such as status, removals, discoveries, and data collection status.
Capacity	Attributes for capacity statistics of a resource, such as available and used capacity, drive capacity, file systems, volumes, and pools, and reserved capacity.

6. To enable the alert for an attribute, click the check mark for the attribute.

7. Specify the conditions for generating an alert for an attribute.

Conditions can include operators such as >=, or <=. Conditions can also include storage values and time values.

For example, for a capacity attribute such as Available Capacity, you can specify that an alert is generated when the amount of available capacity on a resource's pools is less than or equal to 50 GiB.

☒ Available Capacity




Operator	Value	Unit
<=	50	GiB

Tips:

- Not all attributes require criteria for generating alerts. The category and type of an attribute determines whether you can specify criteria and the options that you can select.
- Some attributes can use the operators such as *is*, *is not*, and *contains*. For example, for the Firmware attribute on a DS8000® storage system, select the operator contains. Then, in the value field, type R5 to be notified if the firmware is at the R5 level rather than at a later version such as R6.3, R6.2, or R6.3.

8. Assign a severity to an alert to help you more quickly identify and address the critical conditions that are detected on resources.

The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
 Critical	Assign this severity to alerts that are critical and must be resolved. For example, assign a critical severity to alerts that notify you when the amount of available capacity on a file system falls below a specified threshold.
 Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the status of a data collection job is not normal.
 Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes. For example, assign an informational severity to alerts that are generated when a new fileset is added to a storage system.

9. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the Email Override field.

Tip: If you enter an email address in the Email Override field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to Settings > Notification Settings.



10. Optional: Click View Additional Options to specify how frequently you are notified of alerts.

Use these settings to avoid triggering too many alerts for some conditions.

11. Optional: Click the icon  to duplicate an alert.

Use this action when you want to define another alert for the same attribute but with different criteria and settings.

Duplicating alerts can be helpful in the following situations:

- When you want to generate separate warning alerts and critical alerts for different conditions on the same attribute.
For example, for a capacity attribute such as Available Capacity, you might want to define the following alerts:
 - Define a warning alert  to be generated when the amount of available capacity on a resource's disks is less than or equal to 50 GiB.
 - Duplicate the alert, but this time, specify a critical severity  when the amount of available capacity on a resource's disks is less than or equal to 10 GiB.
- When you want to send alert notifications to different people based on the severity of an alert.
In the previous example for the Available Capacity attribute, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

12. Click Save Changes.

Related reference

- [Alert notifications and actions](#)

Defining general group alerts for performance metrics

You can define alerts that are triggered when the performance of the storage systems or switches that belong to a general group fall outside a specified threshold.

About this task

To define a general group alert for performance metrics, select the metric that you want to measure and specify a threshold value. When the performance of that resource falls outside the threshold, an alert is generated. For example, you can define an alert that notifies you when the back-end response times for managed disks on a SAN Volume Controller exceed 35 milliseconds per read operation. The Overall Back-end Response Time is a metric that measures the average number of milliseconds that it takes to service each read operation on a managed disk.

Procedure

To define a general group alert for performance metrics, complete these steps:

1. In the menu bar, select Groups > General Groups.
2. Right-click a general group in the list and click View Alert Definitions.
3. Click Edit Alert Definitions.
4. Click the type of resource that you want to alert on.
5. Click Performance and then click Add Metrics.
6. Select one or more metrics to alert on and click OK.
7. To enable the alert for a performance metric, click the check mark for the metric.
8. Specify the conditions for generating an alert.

Conditions include an operator and a threshold value.

- a. Select an operator.

An operator determines if an alert is triggered when the performance of a resource is *greater than or equal to* or *less than or equal to* the specified threshold value.

- b. Enter a threshold value.

For example, to trigger an alert if the Total I/O Rate for a storage system is greater than or equal to 500 ops/s, enter the value 500.

Tips for threshold values:

- IBM Spectrum® Control provides recommended values for threshold values that do not vary much between environments. For example, the default threshold values for Port Send Bandwidth Percentage are greater than or equal to 75% for warning alerts, and greater than or equal to 85% for critical alerts.
However, for metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
- For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.

9. Optional: Click View Performance to view a chart of the performance of the resource. Use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.

The chart displays a horizontal color line at the specified threshold value. The color of the line indicates the severity of the alert:




- Critical alert: red
- Warning alert: yellow
- Information alert: blue

For multi-conditional alerts, the chart displays a horizontal line for each condition that shows the threshold value and severity.

To customize the chart, click Top 10 or Bottom 10 to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.

10. Assign a severity to an alert.

Assigning a severity can help you more quickly identify and address the critical performance conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization. Default assignments are provided for each alert.

Option	Description
 Critical	Assign this severity to alerts that are critical and must be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%.
 Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
 Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.


11. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the Email Override field.

Tip: If you enter an email address in the Email Override field, only that email address receives notifications for the alert. The following contacts do not receive notifications:

- Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
- Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to Settings > Notification Settings.



12. Optional: Click View Additional Options to specify how frequently you are notified of alerts.

Use these settings to avoid triggering too many alerts for some conditions.

13. Optional: Click the icon  to duplicate an alert.

Use this action when you want to define another alert for the same metric but with different criteria and settings.

Duplicating alerts can be helpful in the following situations:

- When you want to generate separate warning alerts and critical alerts for different thresholds on the same metric.
For example, for the CRC Error Rate metric for ports, you might want to define the following alerts:
 - Define a warning alert  to be generated when the number of frames per second that are received with cyclic redundancy check (CRC) errors is greater than or equal to 0.01 counts per second.
 - Duplicate the alert, but this time, specify a critical severity  when the CRC error rate is greater than or equal to 0.03 counts per second.
- When you want to send alert notifications to different people based on the severity of an alert.
In the previous example for the CRC Error Rate metric, you can configure the notification settings so that warning alerts are sent to junior administrators, while critical alerts are sent to more senior administrators to resolve.

14. Click Save Changes.

Related tasks

- [Verifying that a performance monitor is running for a resource](#)

Related reference

- [Alert notifications and actions](#)

Defining custom alerts for general groups

Define a custom alert to combine multiple conditions from multiple resources for a general group in a single alert. By creating a custom alert, you can analyze multiple configuration, capacity, and performance conditions together to determine whether an urgent situation occurred on the resources in a group.

About this task




For example, you can create a custom alert that notifies you when the response times for volumes and pools in a group exceed a certain threshold and the performance monitors for storage systems are running without problems.

Tip: In a custom alert, the resource types that you specify conditions for must be associated with the same type of top-level resource. For example, if you include conditions for storage systems, you can also include additional conditions for pools and volumes, but you cannot include conditions for servers.

Procedure

1. In the menu bar, select Groups > General Groups.
2. Right-click a general group in the list and click View Alert Definitions.
3. Click Edit Alert Definitions.
4. Click Custom.
5. Click Create Alert, then enter a name for the custom alert.
6. Assign a severity to the alert.

Assigning a severity can help you more quickly identify and address the critical conditions that are detected on resources. The severity that you assign depends on the guidelines and procedures within your organization.

Option	Description
 Critical	Assign this severity to alerts that are critical and need to be resolved. For example, assign a critical severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 85%. The default severity for custom alerts is critical.
 Warning	Assign this severity to alerts that are not critical, but represent potential problems. For example, assign a warning severity to alerts that notify you when the Port Send Bandwidth Percentage is greater than or equal to 75% but less than 85%.
 Informational	Assign this severity to alerts that might not require any action to resolve and are primarily for informational purposes.

7. Select a component, category, and group for the alert.
For example, select Volumes, Capacity, and Available Capacity.
8. For general and capacity attributes, specify the criteria for generating an alert.
Use criteria such as greater than or equal to, less than or equal to, storage values, and time measurements to customize the conditions under which attributes generate alerts.
For example, for a capacity attribute such as Used Capacity, you can specify that an alert is generated when the amount of used capacity on a resource is more than or equal to 75%. The operator (less than or equal to) + a specified amount of space (75) + the unit of measurement (%) is the criteria that determines if an alert is generated.
Tips:
 - Not all attributes require criteria for generating alerts. The category and type of an attribute determines whether you can specify criteria and the options that you can select.
 - Some attributes can use the operators such as *is*, *is not*, and *changes*. For example, for the Firmware attribute on a DS8000® storage system, select the operator contains. Then, in the value field, type R5 to be notified if the firmware is at the R5 level rather than at a later version such as R6.3, R6.2, or R6.3. This alert definition might be useful if you want to be notified when the firmware for a storage system was reverted to a previous version for some reason.
9. For performance attributes, specify the criteria for generating an alert.
Criteria includes an operator and a threshold value.
 - a. Select an operator.
An operator determines if an alert is triggered when the performance of a resource is *greater than or equal to* or *less than or equal to* the specified threshold value.
 - b. Enter a threshold value.
For example, to trigger an alert if the Total I/O Rate for a volume is greater than or equal to 500 ops/s, enter the value 500.
Tips for threshold values:
 - For metrics that measure throughput and response times, thresholds can vary because of workload, model of hardware, amount of cache memory, and other factors. In these cases, there are no recommended values. To help determine threshold values for a resource, collect performance data over time to establish a baseline of the normal and expected performance behavior for that resource. After you determine a set of baseline values, define alerts to trigger if the measured performance behavior falls outside the normally expected range.
 - For some metrics, lower values might indicate more stress and higher values might indicate idle behavior. For example, a lower threshold value for the Cache Holding Time Threshold metric might indicate a performance problem.
10. Optional: Click View Performance to view a chart of the performance of the resource. Use the chart to evaluate the current and historical performance of a resource to help determine the threshold value for an alert.
The chart uses colored lines to represent the different threshold values and severities that can be defined for an alert:
 - Critical alert: red
 - Warning alert: orange
 - Information alert: blueTo customize the chart, click Top 10 or Bottom 10 to show resources according to their performance, click a time period, and change the start and end dates for the data that is displayed.
11. Optional: If you want to send email notifications of alert violations to contacts other than the policy contacts or global alert notification addresses, enter the email addresses in the Email Override field.
Tip: If you enter an email address in the Email Override field, only that email address receives notifications for the alert. The following contacts do not receive notifications:
 - Any email addresses that are specified as policy contacts, if the alert is in an alert policy.
 - Any global email addresses that are specified for alert notifications. To view the global alert notification addresses, go to Settings > Notification Settings.
12. + Add another condition to the custom alert and repeat steps 7 - 11.
If you do not want to add another condition, continue to the next step.
13. Optional: Click View Additional Options to specify how frequently you are notified of alerts.
Use these settings to avoid triggering too many alerts for some conditions.
14. Click Save Changes.

Related reference

- [Alert notifications and actions](#)

Configuring alert notifications

Alerts can define notification actions that send email, generate Simple Network Management Protocol (SNMP) traps, or generate IBM® Tivoli® Netcool®/OMNIBus events. To enable these notification actions, you must configure IBM Spectrum® Control for email, SNMP, or Tivoli Netcool/OMNIBus alert notifications.

Tip: For alert notifications to generate SNMP traps or IBM Tivoli Netcool/OMNIBus events, first configure IBM Spectrum Control for these notification actions. Then, enable the actions as alert notifications for an alert policy, a resource, or an individual alert.

- [Configuring email alert notifications](#)
To send alert notifications and reports, you must configure the email server.
- [Configuring SNMP alert notifications](#)
You can define an alert to generate SNMP traps when an alert condition is detected on a monitored resource. To enable SNMP alert notifications, configure IBM Spectrum Control to direct the traps to at least one SNMP destination.
- [Configuring Tivoli Netcool/OMNIBus alert notifications](#)
You can define an alert to generate Tivoli Netcool/OMNIBus events when an alert condition is detected on a monitored resource. To enable Tivoli Netcool/OMNIBus notifications, configure IBM Spectrum Control to direct the events to a Tivoli Netcool/OMNIBus server.

Configuring email alert notifications

To send alert notifications and reports, you must configure the email server.

Before you begin

To add or modify settings for the email server, you must be assigned the Administrator role.

About this task

The email server that is used to send alert notifications is also used to send reports. If you didn't configure the email server for alert notifications, you must configure the email server before you create reports.

IBM Spectrum® Control ensures, when it sends alert notifications and reports by email to your email server, that it complies with the security standards that you configure on your email server. For example, if your email server requires authentication and TLS to send emails, then IBM Spectrum Control uses the authentication credentials that you provide when you set up the email server, and uses TLS to establish the connection with your email server. Because the email server that you set up is automatically trusted, you don't have any additional management tasks such as importing the server certificate into the keystores that are used by IBM Spectrum Control.

Procedure

1. Click Settings > Notification Settings.
2. Type the host name, or IPv4, or IPv6 address of the email server.
3. Type the port number that is used by the email server to send alert notifications and reports.
4. Type the name of the user and the password that you want to use to authenticate with the email server.
5. Optional: Type the email address, such as the email address of the administrator, that you want to use for receiving replies to alert notifications and reports.
If you don't enter a value, and you entered an email address in the User Name field, then replies are sent to that address. Otherwise, replies are sent to **no-reply@<Spectrum_Control_server_host_name>**.
Tip: Before you save the settings for the email server, test the connection.
6. Optional: To specify the email addresses that you want to notify when an alert notification is generated, click Global email notification settings, and add each email address separated by a comma.
You can override the global setting for alert notifications by editing the notification settings for a specific alert policy, for a specific resource, or for a specific alert definition that is associated with a resource.
Restriction: Applies only to alert notifications.
7. Click Save.

Configuring SNMP alert notifications

You can define an alert to generate SNMP traps when an alert condition is detected on a monitored resource. To enable SNMP alert notifications, configure IBM Spectrum® Control to direct the traps to at least one SNMP destination.

Before you begin

To modify alert notification settings, you must be assigned the Administrator role.

You must also be running an SNMP management application in your system environment. A Management Information Base (MIB) file is provided in the IBM Spectrum Control installation. You must set up the SNMP trap receiver with the MIB file to receive SNMP traps.

The MIB is automatically deployed when you install IBM Spectrum Control, and can also be found on the installation media.

Table 1. Location of MIB file

Location of MIB file	Directory
Installation media	\data\snmp\tivoliSRM.mib
After IBM Spectrum Control is installed	installation_dir\data\snmp\tivoliSRM.mib

SNMP version support: »IBM Spectrum Control supports SNMPv1, SNMPv2, and SNMPv3 for using SNMP traps destinations to receive alert notifications.⏏

Procedure

To configure SNMP alert notifications, complete the following steps:

1. Go to Settings > Notification Settings.
2. Click the SNMP tab.
3. Click Edit, and specify new settings or remove existing settings.
 - To specify new settings, complete the following steps:
 - Specify the following SNMP alert notification settings for each SNMP destination. You can specify up to two SNMP destinations.

Community

Specify the name of the SNMP community for sending SNMP traps. By default, the community is *public*.

IP address

Specify the host that is configured to receive SNMP traps. You can specify a host name, an IPv4 address, or an IPv6 address depending on what is supported within your environment.

Port

Specify the port number for receiving SNMP traps. SNMP trap messages are sent to this port when an alert condition is detected. By default, the port is set to 162.

- Click Save.
- To delete the current configuration settings without specifying new settings, click Remove.

Configuring Tivoli Netcool/OMNIBus alert notifications

You can define an alert to generate Tivoli® Netcool®/OMNIBus events when an alert condition is detected on a monitored resource. To enable Tivoli Netcool/OMNIBus notifications, configure IBM Spectrum® Control to direct the events to a Tivoli Netcool/OMNIBus server.

Before you begin

To modify alert notification settings, you must be assigned the Administrator role.
In your system environment, you must be running a Tivoli Netcool/OMNIBus server that is configured to receive IBM Spectrum Control alerts.

Procedure

To configure Tivoli Netcool/OMNIBus alert notifications, complete the following steps:

1. Go to Settings > Notification Settings.
2. Click Netcool/OMNIBus in the left-side navigation.
3. Click Edit and specify new settings or remove existing settings.
 - To specify new settings, complete the following steps:
 - Specify the following Tivoli Netcool/OMNIBus alert notification settings:

IP address

Specify the NetCool/OMNIBus server that is configured to receive notifications. You can specify a host name, and IPv4 address, or an IPv6 address. You cannot use a server that is configured for IPv6 only. The server must be IPv4 or dual stack enabled.

Port

Specify the port for receiving alert notifications. Alert notifications are sent to this port when an alert condition is detected. By default, the port is set to 5529.

- Click Save.
 - To delete the current configuration settings without specifying new settings, click Remove.
4. Copy the EIF rules files that are provided with IBM Spectrum Control to the appropriate directory in the Tivoli Netcool/OMNIBus installation.

Rules file names	Location of rules files in IBM Spectrum Control	Copy rules files to this location in Tivoli Netcool/OMNIBus
<ul style="list-style-type: none">• tivoli_eif_tpc.rules• tivoli_eif_tpc_tbsm.rules	<i>extracted_installation_image/data/omnibus</i>	<i>\$OMNIHOME/probes/arch/</i>

For more information about rules file, see the [Tivoli Netcool/OMNIBus Knowledge Center](#).

Triggering conditions for alerts

Define alerts so that IBM Spectrum® Control automatically notifies you when certain conditions or events are detected on monitored resources. Such conditions are the *triggering conditions* for the alert. The specific conditions that can trigger alerts depend on the type of resource that is being monitored.

- [Triggering conditions for storage system alerts](#)
You can set up IBM Spectrum Control so that it examines the attributes, capacity, and performance of a storage system and notifies you when changes or violations are detected.
- [Triggering conditions for hypervisor alerts](#)
You can set up IBM Spectrum Control so that it examines the attributes and capacity of a hypervisor and notifies you when changes are detected.
- [Triggering conditions for switch alerts](#)
You can set up IBM Spectrum Control so that it examines the attributes and performance of a switch and notifies you when changes or violations are detected.
- [Triggering conditions for fabric alerts](#)
You can set up IBM Spectrum Control so that it examines the attributes of a fabric and notifies you when changes are detected.
- [Triggering conditions for server alerts](#)
You can set up IBM Spectrum Control so that it examines the attributes and capacity of a server and notifies you when changes are detected.

Related concepts

- [Alerting](#)

Related tasks

- [Adding resources](#)

Related information

- [Triggering conditions for alerts](#)
- [Triggering actions and notifications for alerts](#)

Triggering conditions for storage system alerts

You can set up IBM Spectrum® Control so that it examines the attributes, capacity, and performance of a storage system and notifies you when changes or violations are detected.

Alerts can notify you of general, capacity, and performance changes and potential issues on the following resources:

- [Performance alert conditions for storage systems](#)
- [Capacity alert conditions for storage systems](#)
- [General alert conditions for storage systems](#)
- [Triggering conditions for storage system internal resource alerts](#)

Important: Not all the attributes upon which you can alert are listed here. To view a complete list of attributes upon which you can alert, go to Settings > Alert Policies. Double-click a default policy for a storage system. Click Edit Alert Definitions on the Alert Definitions tab. View the attributes that are available in the general, capacity, and performance categories. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active. In the tables, default alerts are marked with an asterisk (*).

Tips:

- The type of storage system determines which attributes and performance conditions are available for alerts. For example, triggering conditions for shares are available only for storage systems that are configured for file storage, such as Storwize® V7000 Unified.
- For capacity attributes, you can generate alerts when the amount of storage is greater than, less than, or equal to a specified value. You can also determine the unit of measurement for the attribute, such as KiB, MiB, GiB, or TiB.
- If you are doing tasks where many volumes are being deleted, you might want to temporarily disable alerts that use the Deleted Volume attribute. For example, you might want to disable Deleted Volume alerts temporarily if you are doing maintenance tasks or decommissioning storage.

Performance alert conditions for storage systems

Define alerts that notify you when the performance of a storage system falls outside a specified threshold. In alerts, you can specify conditions based on metrics that measure the performance of volumes, disks, ports, and nodes. By creating alerts with performance conditions, you can be informed about potential bottlenecks in your storage infrastructure.

Examples:

- You can define an alert to be notified when the average number of I/O operations per second for read and write operations on a storage system's volumes is greater than or equal to a specified threshold. Use this alert to be notified when the workload of a volume is high and you might need to balance that load across other volumes to improve overall performance.
- You can define an alert to be notified when the percentage of the average response time that can be attributed to delays from host systems is greater than or equal to a specified threshold. Use this alert to be notified of slow hosts and fabrics that might not be working efficiently.
- You can also define an alert that notifies you when a metric is less than a specified threshold, such as if you want to identify volumes that might be under used.

Tips:

- The type of storage system determines the metrics that can be alerted upon. For a list of the metrics that are available for each type of storage system, see [Performance metrics](#).
- A performance monitor must collect data about a resource before IBM Spectrum Control can determine whether a threshold is violated and an alert is generated for a performance condition.

Best practice: When you set thresholds for performance conditions, try to determine the best value so you can derive the maximum benefit without generating too many false alerts. Because suitable thresholds are highly dependent on the type of workload that is being run, hardware configuration, the number of physical disks, exact model numbers, and other factors, there are no easy or standard default rules.

A recommended approach is to monitor the performance of resources for a number of weeks and by using this historical data, determine reasonable threshold values for each performance condition. After that is done, you can fine-tune the condition settings to minimize the number of false alerts.

Capacity alert conditions for storage systems

Capacity metadata is aggregated and collected by probes. By default, this metadata is collected once every 24 hours.

Table 1. Triggering attributes and conditions for capacity changes on storage systems

Capacity Attributes	Triggering Conditions for Attributes
---------------------	--------------------------------------

Capacity Attributes	Triggering Conditions for Attributes
Adjusted Used Capacity	<p>The amount of capacity that can be used without exceeding the capacity limit. For example, you set a capacity limit of 80% for your storage systems. You want to get an informational alert when the adjusted used capacity exceeds 60% and a critical alert when the adjusted used capacity exceeds 80%. So, you define an informational alert with these parameters:</p> <p>Adjusted Used Capacity \geq 60%</p> <p>And, you define a critical alert with these parameters:</p> <p>Adjusted Used Capacity \geq 80%</p>
Available Capacity (Previously known as Available Pool Space)	<p>The total amount of the space in the pools that is not allocated to the volumes in the pools. To calculate available capacity, the following formula is used:</p> <p>(pool capacity - used capacity)</p> <p>For XIV® systems, pool capacity is the physical capacity of the pools and does not include the provisioned capacity of the pools.</p>
Available Written Capacity (Previously known as Effective Available Capacity)	The total amount of the provisioned capacity in the pools that is not allocated to the volumes in the pools.
Capacity (Previously known as Pool Capacity)	The amount of space in the pools on the storage system that is available for creating volumes.
Capacity-to-Limit	<p>The amount of capacity that is available for storing data before the capacity limit is reached. For example, if you set a capacity limit, you can define a warning alert when the available capacity, in relation to the capacity limit, falls below the value that you specify such as:</p> <p>Capacity-to-Limit \leq 500 GiB</p>
Compression Savings	The estimated percentage of capacity that is saved by using data compression, across all pools on the storage system. The percentage is calculated across all compressed volumes in the pools and does not include the capacity of non-compressed volumes. Inline compression is a software feature that is supported by FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.
Deduplication Savings	The estimated percentage of capacity that is saved by using data deduplication, across all data reduction pools on the storage system. The percentage is calculated across all deduplicated volumes in the pools and does not include the capacity of volumes that are not deduplicated. Available for FlashSystem A9000, FlashSystem A9000R, and resources that run IBM Spectrum Virtualize 8.1.3 or later.
File System Capacity	The total capacity on all of the file systems on the storage system or filer.
Mapped Capacity (Previously known as Assigned Volume Space)	The total volume space in the storage system that is mapped or assigned to host systems, including child pool capacity.
Overprovisioned Capacity (Previously known as Unallocatable Volume Space)	<p>The capacity that cannot be allocated to volumes because the physical capacity of the pools cannot meet the demands for provisioned capacity.</p> <p>IBM Spectrum Control uses the following formula to determine this value: Provisioned Capacity - Capacity</p> <p>Available only for thin-provisioned volumes.</p>
Provisioned Capacity (Previously known as Total Volume Capacity)	The total storage space on all the volumes in pools. For thin-provisioned and compressed volumes, this value includes provisioned capacity. For volumes with parent pools, this value includes child pool capacity.
Raw Capacity (Previously known as Raw Disk Capacity)	The total unformatted disk capacity of a storage system. When this value is calculated, IBM Spectrum Control does not include the capacity of storage system disks that become missing after data collection.
Total Capacity Savings (Previously known as Total Data Reduction Savings)	<p>The estimated percentage of capacity that is saved by using data deduplication, data compression, and thin provisioning.</p> <p>Available for FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.</p>
Reserved Capacity (Previously known as Reserved Pool Space)	<p>The amount of unused capacity in the pool that is reserved for provisioning and optimization tasks.</p> <p>Pool capacity is reserved when a provisioning or optimization task is created, and used when the task is run.</p>
Safeguarded Capacity	The total amount of capacity that is used to store volume backups that are created by the Safeguarded Copy feature in DS8000®.
Shortfall	<p>The difference between the amount of provisioned capacity that is committed to the volumes in the pools and the actual physical space that is available in the pools. As the provisioned capacity is allocated to the thin-provisioned and compressed volumes, the shortfall increases and becomes more critical.</p> <p>This value is determined by the formula, <i>Overprovisioned Capacity ÷ Committed but Unused Capacity</i></p> <p>For example, the physical capacity of the pools is 70 GiB, but 150 GiB of provisioned capacity was committed to the thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB) with only 20 GiB of available pool capacity (70 GiB - 50 GiB). Because only 20 GiB of the pool capacity is available, 80 GiB of the committed capacity cannot be allocated (100 GiB - 20 GiB).</p>
Snapshot Space	The amount of space that is used by all of the snapshots of the file systems that are associated with the IBM Spectrum Scale cluster.
Unreserved Capacity (Previously known as Unreserved Pool Space)	The amount of space in storage system pools that is not allocated for volumes, and is not reserved by pending or scheduled provisioning tasks.
Unmapped Capacity	The total volume space in the storage system that is not mapped or assigned to host systems.
Unused Volume Capacity (Previously known as Effective Unallocated Volume Space)	The amount of the provisioned capacity in the storage pool that is not used.

Capacity Attributes	Triggering Conditions for Attributes
Used Capacity (Previously known as Physical Allocation)	The percentage of physical capacity in pools that is allocated to volumes, including child pools. The value is always less than or equal to 100% because you cannot allocate more physical capacity to the volumes than is available in the pools. This value is determined by the formula, $Used\ Capacity \div Capacity \times 100$. For example, if the capacity that is reserved for volumes is 50 GiB for a volume size of 200 GiB, used capacity is 25%.
Used Capacity (Previously known as Used Pool Space)	The capacity in the pool that is allocated to and used by volumes.
Used Written Capacity (%) (Previously known as Effective Used Capacity)	The percentage of capacity that is provisioned to the standard-provisioned volumes and the thin-provisioned volume, given the drive compression savings.
Used Written Capacity (Previously known as Effective Used Capacity)	The total amount of provisioned capacity that is used by all volumes, given the drive compression savings.
Written Capacity Limit (Previously known as Effective Capacity)	The amount of provisioned capacity that can be created, given the drive compression savings.

General alert conditions for storage systems

Asset, capacity, and configuration metadata is aggregated and collected when probes collect storage system metadata. By default, metadata is collected once every 24 hours.

Table 2. Triggering attributes and conditions for general changes on storage systems

General Attributes	Triggering Conditions for Attributes
Firmware	<p>The firmware version of the microcode on a storage system. For the DS-series of storage systems, this value represents the SEA version of the firmware.</p> <p>To view information about the code bundles for the firmware versions of the DS-series, go to IBM® Support and search for code bundle information. An internet connection is required to access the support site.</p>
Last Successful Probe	A specified amount of time has passed since a probe or performance monitor was able to collect data about a storage system. You can use this alert to be notified when up-to-date configuration, status, or performance data is not being collected about a storage system and its existing data might be stale. This situation might occur if the resource, network, or IBM Spectrum Control server is unavailable.
Last Successful Monitor	
Performance Monitor Status*	<p>One of the following statuses is detected for a performance monitor:</p> <p>Not Normal An error or warning occurred during data collection. This status indicates that a performance monitor did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A performance monitor completed, but did not collect a complete set of performance data. This status might occur if the resource was rebooted during data collection, no valid performance data was provided by the resource, or a communication error occurred with the resource or its associated agent.</p> <p>Error A performance did not complete when it attempted to collect performance data about the resource. This status might occur if the resource cannot be reached during data collection, or if no configuration data is available for the resource.</p> <p>For details about why a specific status occurred, check the log for the performance monitor. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Performance Monitor section on the Data Collection page.</p>
Probe Status*	<p>One of the following statuses is detected for a probe:</p> <p>Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.</p> <p>Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.</p> <p>For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Probe section on the Data Collection page.</p>
Status	<p>One of the following statuses is detected for a storage system:</p> <p>Not Normal An error or warning status was detected for a storage system.</p> <p>Warning A warning status was detected for a storage system. This status might occur if a storage system comes online or if its version changes.</p> <p>Error An error status was detected for a storage system. This status might occur if the cooling fans in a storage system are stopped and the internal temperature is too high or if a storage system goes offline.</p>

- [Triggering conditions for storage system internal resource alerts](#)

You can set up IBM Spectrum Control so that it examines the attributes, capacity, and performance of the internal resources of storage systems and notifies you when changes or violations are detected.

Triggering conditions for storage system internal resource alerts

You can set up IBM Spectrum® Control so that it examines the attributes, capacity, and performance of the internal resources of storage systems and notifies you when changes or violations are detected.

Alerts can notify you of general changes and performance issues on the following internal resources of storage systems:

- [Internal resources \(common conditions\)](#)
- [Clusters](#)
- [Drives](#)
- [Disk Groups](#)
- [FC Ports](#)
- [Filesets](#)
- [File Systems](#)
- [File System Pools](#)
- [Host Connections](#)
- [I/O Groups](#)
- [IP Ports](#)
- [Managed Disks](#)
- [Modules](#)
- [Network Shared Disks](#)
- [Nodes](#)
- [Pools](#)
- [Quotas](#)
- [RAID Arrays](#)
- [Shares](#)
- [Volumes](#)

Important: Not all the attributes upon which you can alert are listed here. To view a complete list of attributes upon which you can alert, go to Settings > Alert Policies. Double-click a default policy for a storage system. Click Edit Alert Definitions on the Alert Definitions tab. View the attributes that are available in the general, capacity, and performance categories. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active. In the tables, default alerts are marked with an asterisk (*).

Tips:

- The type of storage system determines which attributes and performance conditions are available for alerts. For example, triggering conditions for shares are available only for storage systems that are configured for file storage, such as Storwize® V7000 Unified.
- For capacity attributes, you can generate alerts when the amount of storage is greater than, less than, or equal to a specified value. You can also determine the unit of measurement for the attribute, such as KiB, MiB, GiB, or TiB.

Internal resources (common conditions)

There are a number of alert conditions that are common to many of the internal resources in a storage system. These common conditions represent key changes in your storage infrastructure. For example, you can specify conditions that generate alerts when specific internal resources are added to or deleted from a storage system, or when current data isn't being collected about resources.

By default, asset, capacity, and configuration metadata for storage systems is aggregated and collected daily. Define alerts to track daily changes to the attributes and conditions of the internal resources in your storage systems.

Table 1. Triggering attributes and conditions that are common to internal resources

General Attributes	Triggering Conditions for Attributes
New resource	A resource is detected for the first time. Use this alert to be notified when new physical and logical resources are added to a storage system.
Removed physical resource	A previously monitored resource can no longer be found. Historical data about the resource is retained, but no current data is being collected. Use these alerts to be notified if a physical or logical resource is removed, deleted, or becomes unavailable.
Deleted logical resource	Physical resources include disks, RAID arrays, I/O groups, ports, nodes, host connections, clusters, and file systems. Logical resources include volumes, pools, filesets, and shares.

Clusters

Table 2. Triggering attributes and conditions for general changes on clusters

General Attributes	Triggering Conditions for Attributes
New Cluster	A cluster is detected for the first time.
Removed NAS Cluster	A previously monitored NAS cluster can no longer be found. Historical data about the cluster is retained, but no current data is being collected. Use this alert to be notified if a cluster is removed or becomes unavailable.

Drives

Table 3. Triggering attributes and conditions for general changes on storage system drives

General Attributes	Triggering Conditions for Attributes
New drive	A drive is detected for the first time.
Removed drive	A previously monitored drive can no longer be found. Historical data about the drive is retained, but no current data is being collected. Use this alert to be notified if a drive is removed or becomes unavailable.

Table 4. Triggering attributes and conditions for capacity changes on drives

Capacity Attributes	Triggering Conditions for Attributes
Available Drive Capacity (Previously known as Available Disk Space)	The capacity that is available (not allocated) on the drive.
Capacity (Previously known as Disk Capacity)	The total amount of storage capacity that is on the drive.

Disk Groups

Table 5. Triggering attributes and conditions for general changes on storage system disk groups

General Attributes	Triggering Conditions for Attributes
New Disk Group	A disk group is detected for the first time.
Removed Disk Group	A previously monitored disk group can no longer be found. Historical data about the disk group is retained, but no current data is being collected. Use this alert to be notified if a disk group is removed or becomes unavailable.
Status	<p>One of the following statuses is detected for a disk group:</p> <p>Not Normal An error or warning condition is detected on a disk group.</p> <p>Warning A warning condition is detected on a disk group. This condition might occur if any of the disks in the disk group goes offline.</p> <p>Error An error condition is detected on a disk group. This condition might occur under the following conditions:</p> <ul style="list-style-type: none"> The percentage of remaining, unused volume capacity in a disk group that is not available to be used is too high. The available capacity in a disk group that is not reserved for volumes is too low. The disk group provisioned capacity exceeds the warning or critical threshold boundary value.

FC Ports

Table 6. Triggering attributes and conditions for general changes on Fibre Channel ports

General Attributes	Triggering Conditions for Attributes
New FC Port	A new FC port was detected for the first time.
Removed FC Port	A previously monitored FC port can no longer be found. Historical data about the port is retained, but no current data is being collected. Use this alert to be notified if a port is deleted or becomes unavailable.
Status	<p>One of the following statuses is detected for an FC port:</p> <p>Not Normal An error or warning condition is detected on a port.</p> <p>Warning A warning condition is detected on a port.</p> <p>Error An error condition is detected on a port.</p>

Filesets

Table 7. Triggering attributes and conditions for general changes on filesets

General Attributes	Triggering Conditions for Attributes
Deleted NAS Fileset	A previously monitored NAS fileset can no longer be found. Historical data about the fileset is retained, but no current data is being collected. Use this alert to be notified if a fileset is deleted or becomes unavailable.
New Fileset	A fileset is detected for the first time.
State	A fileset is linked to a file system, unlinked from a file system, or deleting.

File Systems

Table 8. Triggering attributes and conditions for general changes on file systems

General Attributes	Triggering Conditions for Attributes
New File System	A file system is detected for the first time.
Removed NAS File System	A previously monitored NAS file system can no longer be found. Historical data about the file system is retained, but no current data is being collected. Use this alert to be notified if a file system is deleted or becomes unavailable.

File System Pools

Table 9. Triggering attributes and conditions for general changes on file system pools

General Attributes	Triggering Conditions for Attributes
Deleted Pool	A previously monitored NAS pool can no longer be found. Historical data about the pool is retained, but no current data is being collected. Use this alert to be notified if a pool is deleted or becomes unavailable.
New Pool	A NAS pool is detected for the first time.

Host Connections

Table 10. Triggering attributes and conditions for general changes on host connections

General Attributes	Triggering Conditions for Attributes
New Host Connection	A host connection is detected for the first time.
Removed Host Connection	A previously monitored host connection can no longer be found. Historical data about the host connection is retained, but no current data is being collected. Use this alert to be notified if a host connection is removed or becomes unavailable.

Table 11. Triggering attributes and conditions for changes to performance metrics for unmap operations on host connections

Performance Attribute	Triggering Conditions for Attributes
Data Rate (Unmap)	Define an alert to monitor the average number of MiBs per second that were unmapped. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Overall I/O Rate (Unmap)	Define an alert to monitor the average number of unmap operations per second. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Peak Response Time (Unmap)	Define an alert to monitor the worst response time measured for an unmap operation in the sample interval. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Response Time (Unmap)	Define an alert to monitor the average number of milliseconds required to complete an unmap operation. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Unaligned Unmap I/O Rate	Define an alert to monitor the average number of volumes unmap operations per second that are not aligned on an 8K boundary. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.

I/O Groups

Table 12. Triggering attributes and conditions for general changes on I/O groups

General Attributes	Triggering Conditions for Attributes
New I/O Group	A new I/O group was detected for the first time.
Removed I/O Group	A previously monitored I/O group can no longer be found. Historical data about the I/O group is retained, but no current data is being collected. Use this alert to be notified if an I/O group is removed or becomes unavailable.

IP Ports

Triggering attributes and conditions are available for Internet Protocol ports on the nodes on block storage systems that run IBM Spectrum Virtualize, such as SAN Volume Controller, the IBM® Storwize family, and some models of the IBM FlashSystem® family.

Table 13. Triggering attributes and conditions for general changes on IP ports

General Attributes	Triggering Conditions for Attributes
Status	One of the following statuses was detected for an IP port: <ul style="list-style-type: none"> Configured Unconfigured
Removed Port	A previously monitored IP port can no longer be found. Historical data about the port is retained, but no current data is being collected. Use this alert to be notified if a port is deleted or becomes unavailable.
New Port	A new IP port was detected for the first time.
Host Attach	An IP port was attached to, or detached from, a host.
Storage Attach	An IP port was attached to, or detached from, a storage system.
Management	One of the following management statuses was detected for an IP port: <ul style="list-style-type: none"> Configured Unconfigured
Remote Copy Relationship	The remote copy relationship changed, or one of the following remote copy statuses was detected for an IP port: <ul style="list-style-type: none"> Active Unconfigured Standby

Managed Disks

Table 14. Triggering attributes and conditions for general changes on managed disks

General Attributes	Triggering Conditions for Attributes
Managed Disk Status	One of the following statuses is detected for a managed disk: Not Normal An error or warning status was detected on a managed disk. Warning A warning status was detected on a managed disk. Error An error status was detected on a managed disk.
New Managed Disk	A new managed disk was detected for the first time.
Removed Managed Disk	A previously monitored managed disk can no longer be found. Historical data about the managed disk is retained, but no current data is being collected. Use this alert to be notified if a managed disk is removed or becomes unavailable.

Table 15. Triggering attributes and conditions for capacity changes on managed disks

Capacity Attributes	Triggering Conditions for Attributes
Available Capacity	The unused storage capacity on the managed disk.
Capacity (Previously known as Total Space)	The total capacity on the managed disk on the storage system. This attribute is only available for Storwize V7000 storage systems that are configured as back-end storage.

Modules

Table 16. Triggering attributes and conditions for general changes on modules

General Attributes	Triggering Conditions for Attributes
Deleted Module	A previously monitored module can no longer be found. Use this alert to be notified if a module is removed or becomes unavailable.

Network Shared Disks

Table 17. Triggering attributes and conditions for general changes on NSDs

General Attributes	Triggering Conditions for Attributes
New NSD	A new NSD was detected for the first time.
Removed NSD	A previously monitored NSD can no longer be found. Historical data about the NSD is retained, but no current data is being collected. Use this alert to be notified if an NSD is removed or becomes unavailable.
Status	One of the following statuses is detected for an NSD: Not Normal An error or warning status was detected on an NSD. Warning A warning status was detected on an NSD. This status might occur if a storage system comes online or if its version changes. Error An error status was detected on an NSD.

Nodes

Table 18. Triggering attributes and conditions for general changes on nodes

General Attributes	Triggering Conditions for Attributes
New Node	A new node was detected for the first time.
Removed Node (Block storage)	A previously monitored node can no longer be found. Historical data about the node is retained, but no current data is being collected. Use this alert to be notified if a node is removed or becomes unavailable.
Removed NAS Node (File storage)	A previously monitored NAS node can no longer be found. Historical data about the node is retained, but no current data is being collected. Use this alert to be notified if a node is removed or becomes unavailable.
Cloud Gateway Status (File storage)	One of the following statuses is detected for a node. This attribute is available only for IBM Spectrum Scale. Not Running The gateway service is not running because, for example, the service is stopped, no cloud account is configured, or the connection to the cloud provider failed. Not Installed The node is not a cloud gateway. No File System The node is a cloud gateway and the gateway service is running, but the node is not yet assigned to a file system. Stopped The node is a cloud gateway for a file system, but the gateway service is stopped. No Cloud Account The node is a cloud gateway for a file system and the gateway service is running, but no cloud account is configured. Disconnected The node is a cloud gateway, the gateway service is running, and a cloud account is configured, but the connection to the cloud provider failed.

Table 19. Triggering attributes and conditions for changes to performance metrics for cache fullness on nodes

Performance Attributes	Triggering Conditions for Attributes
Max Read Cache Fullness	Define an alert to monitor the maximum amount of the lower cache which the cache partitions of the pools that are managed by the node use for read operations. If the maximum value for the cache reaches 100%, the read cache partition for one or more of the pools is full. The read operations that pass through the node to the affected pools will be queued and the I/O response times will increase for the volumes in the affected pools. This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.
Max Write Cache Fullness	Define an alert to monitor the maximum amount of the lower cache which the cache partitions of the pools that are managed by the node use for write operations. If the maximum value for the cache reaches 100%, the write cache partition for one or more of the pools is full. The write operations that pass through the node to the affected pools will be queued and the I/O response times will increase for the volumes in the affected pools. This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.
Read Cache Fullness	Define an alert to monitor the average amount of the lower cache which the cache partitions of the pools that are managed by the node use for read operations. Use this alert to monitor the average cache fullness for read operations to identify the nodes that experience heavy cache usage. This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.
Write Cache Fullness	Define an alert to monitor the average amount of the lower cache which the cache partitions of the pools that are managed by the node are using for write operations. Use this alert to monitor the average cache fullness for write operations to identify the nodes that experience heavy cache usage. This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.

Table 20. Triggering attributes and conditions for changes to performance metrics for recovering data in data reduction pools on nodes

Performance Attributes	Triggering Conditions for Attributes
Data Rewrite Rate	Define an alert to monitor the rate at which data is rewritten when a host overwrites data in data reduction pools on the node. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.
Extent Collection Rate	Define an alert to monitor the number of volume extents that were processed for garbage collection. The reclaimable capacity in the volume extents is collected so that it can be reused in the data reduction pools on the node. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.
Data Movement Rate	Define an alert to monitor the rate at which valid data in a reclaimed volume extent is moved to a new extent in the data reduction pool on the node. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.
New Address Write Rate	Define an alert to monitor the rate at which capacity is used to write the host's data to unallocated addresses in the data reduction pool on the node. Use this alert to determine which hosts are increasing the amount of capacity that is being written to data reduction pools on a node. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.
Reclaimable Capacity	Define an alert to monitor the amount of capacity that can be reclaimed in the data reduction pools on the node. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.
Recovered Capacity Rate	Define an alert to monitor the rate at which capacity is recovered by garbage collection for reuse in the data reduction pools on the node. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.2 or later.

Pools

Table 21. Triggering attributes and conditions for general changes on pools

General Attributes	Triggering Conditions for Attributes
Deleted Pool	A previously monitored pool can no longer be found. Historical data about the pool is retained, but no current data is being collected. Use this alert to be notified if a pool is deleted or becomes unavailable.
New Storage Pool	A new pool was detected for the first time.
Status	One of the following statuses is detected for a pool: Not Normal An error or warning condition is detected on a pool. Warning A warning condition is detected on a pool. This condition might occur if a pool goes offline. Error An error condition is detected on a pool. This condition might occur under the following conditions: <ul style="list-style-type: none"> The percentage of remaining, unused volume capacity in a pool that is not available to be used is too high. The available capacity in a pool that is not reserved for volumes is too low. The pool total provisioned capacity exceeds the warning or critical threshold boundary value.

Table 22. Triggering attributes and conditions for capacity changes on pools

Capacity Attributes	Triggering Conditions for Attributes
Adjusted Used Capacity	The amount of capacity that can be used without exceeding the capacity limit. For example, you set a capacity limit of 80% for your pools. You want to get an informational alert when the adjusted used capacity exceeds 60% and a critical alert when the adjusted used capacity exceeds 80%. So, you define an informational alert with these parameters: Adjusted Used Capacity ≥ 60% And, you define a critical alert with these parameters: Adjusted Used Capacity ≥ 80%

Capacity Attributes	Triggering Conditions for Attributes
Available Repository Capacity	The amount of available, unallocated storage space on all extents in the repository of a pool for Track Space-Efficient (TSE) thin-provisioning. This attribute applies only to the DS8000® storage systems. You can use this alert to be notified about space-efficient volumes.
Available Virtual Capacity	The amount of provisioned capacity in a thin-provisioned pool that is not used by volumes.
Available Written Capacity (Previously known as Effective Available Capacity)	The total amount of the provisioned capacity in the pools that is not allocated to the volumes in the pools.
Capacity-to-Limit	The amount of capacity that is available for storing data before the capacity limit is reached. For example, if you set a capacity limit, you can define a warning alert when the available capacity relative to the capacity limit, falls below the value that you specify, such as: Capacity-to-Limit ≤ 500 GiB
Compression Savings	The estimated percentage of capacity that is saved by using data compression. The percentage is calculated across all compressed volumes in the pool and does not include the capacity of non-compressed volumes. Inline compression is a software feature that is supported by FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV® storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.
Deduplication Savings	The estimated percentage of capacity that is saved by using data deduplication. The percentage is calculated across all deduplicated volumes in the pool and does not include the capacity of volumes that are not deduplicated. Available for resources that run IBM Spectrum Virtualize 8.1.3 or later.
Mapped Capacity (Previously known as Assigned Volume Space)	The amount of space on all the volumes in a pool that are mapped or assigned to host systems. For a thin-provisioning pool, this value includes the provisioned capacity of thin-provisioned volumes, which might exceed the total capacity in the pool. For Hitachi VSP non-thin provisioning pool capacity, this value is the sum of assigned regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.
Enterprise HDD Capacity	The total capacity on Enterprise hard disk drives in the pool. Easy Tier® can use the capacity on these drives to retire the volume extents in the pool.
Enterprise HDD Available Capacity	The available capacity on Enterprise hard disk drives in the pool. Easy Tier can use the available capacity to retire the volume extents in the pool.
Enterprise HDD Available Capacity (%)	The percentage of capacity on Enterprise hard disk drives in the pool that is available. Easy Tier can use the available capacity to retire the volume extents in the pool.
Nearline HDD Capacity	The total capacity on Nearline hard disk drives in the pool. Easy Tier can use the capacity on these drives to retire the volume extents in the pool.
Nearline HDD Available Capacity	The available capacity on Nearline hard disk drives in the pool. Easy Tier can use the available capacity to retire the volume extents in the pool.
Nearline HDD Available Capacity (%)	The percentage of capacity on Nearline hard disk drives in the pool that is available. Easy Tier can use the available capacity to retire the volume extents in the pool.
Provisioned Capacity (Previously known as Virtual Allocation)	The percentage of physical space in a pool that is committed to the provisioned capacity of the volumes in the pool. You can use this alert to be notified about space-efficient volumes.
Repository Space	The amount of space on all extents in the repository of a pool. This attribute applies only to the DS8000 storage systems.
Reserved Capacity (Previously known as Reserved Pool Space)	The amount of unused capacity in a pool that is reserved for provisioning and optimization tasks. Pool capacity is reserved when a provisioning or optimization task is created, and used when the task is run.
Reserved Volume Capacity (Previously known as Unused Space)	The amount of pool capacity that is reserved but has not been used yet to store data on the thin-provisioned volume.
Safeguarded Capacity	The total amount of capacity that is used to store volume backups that are created by the Safeguarded Copy feature in DS8000.
Shortfall	The difference between the amount of provisioned capacity that is committed to the volumes in the pools and the actual physical capacity that is available in the pools. As the provisioned capacity is allocated to the thin-provisioned and compressed volumes, the shortfall increases and becomes more critical. This value is determined by the formula, <i>Overprovisioned Capacity ÷ Committed but Unused Capacity</i> For example, the physical capacity of the pools is 70 GiB, but 150 GiB of provisioned capacity was committed to the thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB – 50 GiB) with only 20 GiB of available pool capacity (70 GiB – 50 GiB). Because only 20 GiB of the pool capacity is available, 80 GiB of the committed capacity cannot be allocated (100 GiB – 20 GiB).
SCM Capacity	The total capacity on Storage Class Memory (SCM) drives in the pool. Easy Tier can use the capacity on SCM drives to retire the volume extents in the pool. Available for IBM Spectrum Virtualize systems, such as IBM FlashSystem 9100, IBM FlashSystem 7200, and the IBM Storwize family.
SCM Available Capacity	The available capacity on Storage Class Memory (SCM) drives in the pool. Easy Tier can use the available capacity to retire the volume extents in the pool. Available for IBM Spectrum Virtualize systems, such as IBM FlashSystem 9100, IBM FlashSystem 7200, and the IBM Storwize family.
SCM Available Capacity (%)	The percentage of capacity on Storage Class Memory (SCM) drives in the pool that is available. Easy Tier can use the available capacity to retire the volume extents in the pool. Available for IBM Spectrum Virtualize systems, such as IBM FlashSystem 9100, IBM FlashSystem 7200, and the IBM Storwize family.
Tier 0 Flash Capacity	The total capacity on Tier 0 flash solid-state drives in the pool. Easy Tier can use the capacity on these drives to retire the volume extents in the pool.
Tier 0 Flash Available Capacity	The available capacity on Tier 0 flash solid-state drives in the pool. Easy Tier can use the available capacity to retire the volume extents in the pool.
Tier 0 Flash Available Capacity (%)	The percentage of capacity on Tier 0 flash solid-state drives in the pool that is available. Easy Tier can use the available capacity to retire the volume extents in the pool.

Capacity Attributes	Triggering Conditions for Attributes
Tier 1 Flash Capacity	The total capacity on Tier 1 flash solid-state drives in the pool. Easy Tier can use the capacity on these drives to retier the volume extents in the pool.
Tier 1 Flash Available Capacity	The available capacity on Tier 1 flash solid-state drives in the pool. Easy Tier can use the available capacity to retier the volume extents in the pool.
Tier 1 Flash Available Capacity (%)	The percentage of capacity on Tier 1 flash solid-state drives in the pool that is available. Easy Tier can use the available capacity to retier the volume extents in the pool.
Tier 2 Flash Capacity	The total capacity on Tier 2 flash, high-capacity drives in the pool. Easy Tier can use the capacity on these drives to retier the volume extents in the pool. Available for DS8000 storage systems.
Tier 2 Flash Available Capacity	The available capacity on Tier 2 flash, high-capacity drives in the pool. Easy Tier can use the available capacity to retier the volume extents in the pool. Available for DS8000 storage systems.
Tier 2 Flash Available Capacity (%)	The percentage of capacity on Tier 2 flash, high-capacity drives in the pool that is available. Easy Tier can use the available capacity to retier the volume extents in the pool. Available for DS8000 storage systems.
Total Capacity Savings (Previously known as Total Data Reduction Savings)	The estimated percentage of capacity that is saved by using data deduplication, data compression, and thin provisioning. Available for FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.
Total Reserved Capacity	The total amount of space on the pool that is reserved for provisioning and optimization tasks. Pool space is reserved when a provisioning or optimization task is created, and allocated when the task is run.
Used Capacity (Previously known as Allocated Space)	The amount of space that is reserved for all the volumes in a pool.
Used Capacity (%) (Previously known as Physical Allocation)	The percentage of physical capacity in the pool that is used by the standard-provisioned volumes, the thin-provisioned volumes, and the volumes in child pools. This value is always less than or equal to 100% because you cannot allocate more physical space than is available in a pool. This value is determined by the formula, $Used\ Capacity \div Capacity \times 100$. For example, if the space that is reserved for volumes is 50 GiB for a volume size of 200 GiB, used capacity is 25%.
Used Repository Space	The amount of used capacity on all extents in the repository of a pool. This attribute applies only to the DS8000 storage systems. You can use this alert to be notified about space-efficient volumes.
Used Volume Space	The amount of space on the storage system that is used by volumes.
Used Written Capacity (%) (Previously known as Effective Used Capacity)	The percentage of capacity that is provisioned to the standard-provisioned volumes and the thin-provisioned volumes, given the drive compression savings.
Used Written Capacity (Previously known as Effective Used Capacity)	The total amount of provisioned capacity that is used by all the volumes given the drive compression savings.
User Reserved Capacity	The amount of space in the pools on the storage system that is reserved for user-defined purposes.
Virtual Capacity Limit	The maximum amount of virtual storage space available to allocate to volumes in the storage pools that are associated with the storage system. You can use this alert to be notified about space-efficient volumes.
Virtual Volume Space	The total amount of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. You can use this alert to be notified about space-efficient volumes.
Written Capacity Limit (Previously known as Effective Capacity)	The amount of provisioned capacity that can be created, given the drive compression savings.

Table 23. Triggering attributes and conditions for changes to performance metrics for cache fullness on pools

Performance Attributes	Triggering Conditions for Attributes
Read and Write Cache Fullness	Define an alert to monitor the average amount of the lower cache which the pools' cache partitions on the nodes use for read and write operations. Use this alert to monitor the average cache fullness for read and write operations to identify the pools that experience heavy cache usage. This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.
Max Read and Write Cache Fullness	Define an alert to monitor the maximum amount of the lower cache, which the cache partitions on the nodes that manage the pool use for read and write operations. If the maximum value for the cache reaches 100%, one or more cache partitions on one or more pools is full. The operations that pass through the pools with full cache partitions will be queued and I/O response times will increase for the volumes in the affected pools. This metric applies to systems that are running IBM Spectrum Virtualize V7.3 or later.

Quotas

Table 24. Triggering attributes and conditions for general changes on quotas

General Attributes	Triggering Conditions for Attributes
Deleted Quota	A previously monitored quota can no longer be found. Historical data about the quota is retained, but no current data is being collected. Use this alert to be notified if a quota is deleted or becomes unavailable.
New Quota	A new quota was detected for the first time.

RAID Arrays

Table 25. Triggering attributes and conditions for general changes on RAID arrays

General Attributes	Triggering Conditions for Attributes
New RAID Array	A new RAID array was detected for the first time.
Removed RAID Array	A previously monitored RAID array can no longer be found. Historical data about the RAID array is retained, but no current data is being collected. Use this alert to be notified if a RAID array is removed or becomes unavailable.
Status	One of the following statuses is detected for a RAID array: Not Normal An error or warning condition is detected on a RAID array. Warning A warning condition is detected on a RAID array. Error An error condition is detected on a RAID array.

Table 26. Triggering attributes and conditions for capacity changes on RAID arrays

Capacity Attributes	Triggering Conditions for Attributes
Available Physical Capacity, Available Physical Capacity (%)	The amount and percentage of storage space that is unused on all the disk drive modules (DDMs) in the RAID array. Available for RAID arrays with disk drive modules that use inline data compression, such as RAID arrays on FlashSystem 9100 and FlashSystem 900.
Compression Savings, Compression Savings (%)	For compressed RAID arrays, the amount and percentage of capacity that is saved by using drive compression.
Capacity (Previously known as Total Space)	For uncompressed RAID arrays, the total capacity is the same as the physical capacity and represents the total storage capacity of all the DDMs in the array. For compressed RAID arrays, the total capacity is the estimated amount of data that can be written to the array. This value is larger than the physical capacity as the drive compression is used to reduce the size of the data.

Table 27. Triggering attributes and conditions for changes to performance metrics for RAID arrays

Performance Attribute	Triggering Conditions for Attributes
Data Rate	The average rate at which data is transferred in MiB per second.
I/O Rate	The average number of operations per second.
Response Time	The average number of milliseconds required to complete an operation.
Transfer size	The average number of KiB that are transferred per I/O operation.
Disk Utilization Percentage	The average percentage of time the disks that are associated with an array are busy.
Sequential I/O Percentage	The percentage of operations among all I/O operations that were sequential I/O operations.

Shares

Table 28. Triggering attributes and conditions for general changes on shares

General Attributes	Triggering Conditions for Attributes
New Share	A new share was detected for the first time.
Deleted Share	A previously monitored share can no longer be found. Historical data about the share is retained, but no current data is being collected. Use this alert to be notified if a share is deleted or becomes unavailable.

Volumes

Table 29. Triggering attributes and conditions for general changes on volumes

General Attributes	Triggering Conditions for Attributes
Deleted Volume	A previously monitored volume can no longer be found. Historical data about the volume is retained, but no current data is being collected. Use this alert to be notified if a volume is deleted or becomes unavailable. If you are doing tasks where many volumes are being deleted, you might want to temporarily disable alerts that use the Deleted Volume attribute. For example, you might want to disable Deleted Volume alerts temporarily if you are doing maintenance tasks or decommissioning storage.
New Volume	A new volume was detected for the first time.
Status	One of the following statuses is detected for a volume: Not Normal An error or warning condition is detected on a RAID array. Warning A warning condition is detected on a RAID array. Error An error condition is detected on a RAID array.

Table 30. Triggering attributes and conditions for capacity changes on volumes

Capacity Attributes	Triggering Conditions for Attributes
Compression Savings	The estimated percentage of capacity that is saved by using data compression. Inline compression is a software feature that is supported by FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.

Capacity Attributes	Triggering Conditions for Attributes
Provisioned Capacity (Previously known as Unallocatable Space)	The amount of space by which the capacity of a volume exceeds the physical capacity of the associated pool. In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool by creating volumes with more virtual capacity than can be physically allocated in the pool. This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.
Real Capacity	The total amount of storage space that is physically allocated to a volume. For thin-provisioned volumes, this value is less than the provisioned capacity of the volume. In an XIV and IBM Spectrum Accelerate, this value represents the physical ("hard") capacity of the volume, not the provisioned ("soft") capacity.
Remaining Managed Space	The amount of storage space that is available on a managed disk. This value is only available for Storwize V7000 storage systems that are configured as back-end storage.
Reserved Volume Capacity (Previously known as Unused Space)	The amount of pool capacity that is reserved but has not been used yet to store data on the thin-provisioned volume. The value for Reserved Volume Capacity is available only for SAN Volume Controller and Storwize family storage systems that are configured with block storage.
Safeguarded Capacity	The amount of capacity that is used to store volume backups that are created by the Safeguarded Copy feature in DS8000.
Space	The amount of space in a pool that is allocated to a volume.
Unused Capacity (Previously known as Unallocated Space)	The capacity in a pool that is not reserved for a volume. This value is determined by the formula: <i>Capacity – Used Capacity</i> The value for Unused Capacity is available only for thin provisioned volumes.
Uncompressed Used Capacity	The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.
Used Capacity (Previously known as Allocated Space)	The capacity on a pool that is physically allocated to a volume.
Written Capacity (Previously known as Written Space)	The amount of data that is written from the assigned hosts to the volume before compression or data deduplication are used to reduce the size of the data.

Table 31. Triggering attributes and conditions for changes to performance metrics for unmap operations on volumes

Performance Attribute	Triggering Conditions for Attributes
Data Rate (Unmap)	Define an alert to monitor the average number of MiBs per second that were unmapped. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Overall I/O Rate (Unmap)	Define an alert to monitor the average number of unmap operations per second. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Peak Response Time (Unmap)	Define an alert to monitor the worst response time measured for an unmap operation in the sample interval. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Response Time (Unmap)	Define an alert to monitor the average number of milliseconds required to complete an unmap operation. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.
Unaligned Unmap I/O Rate	Define an alert to monitor the average number of volumes unmap operations per second that are not aligned on an 8K boundary. This metric applies to systems that are running IBM Spectrum Virtualize V8.1.1 or later.

Triggering conditions for hypervisor alerts

You can set up IBM Spectrum® Control so that it examines the attributes and capacity of a hypervisor and notifies you when changes are detected.

Alerts can notify you of general changes and capacity changes on the following resources:

- [Hypervisors](#)
- [Controllers](#)
- [Data Stores](#)
- [Disks](#)
- [Paths](#)
- [Virtual Machines](#)
- [VMDKs](#)

Important: Not all the attributes upon which you can alert are listed here. A number of other attributes are available for alerts and are based on the key properties of a hypervisor. To view a complete list of hypervisor attributes upon which you can alert, go to [Settings > Alert Policies](#). Double-click a default policy for a hypervisor. Click Edit Alert Definitions on the Alert Definitions tab. View the attributes that are available in the general and capacity categories. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active.

In the tables, default alerts are marked with an asterisk (*).

Tip: For capacity attributes, you can generate alerts when the amount of storage is greater than, less than, or equal to a specified value. You can also determine the unit of measurement for the attribute, such as KiB, MiB, GiB, or TiB.

Hypervisors

Table 1. Triggering attributes and conditions for general hypervisor changes

General Attributes	Descriptions of Triggering Conditions
Last Successful Probe	A specified amount of time has passed since a probe was able to collect data about a hypervisor. You can use this alert to be notified when up-to-date configuration and status data is not being collected about a hypervisor and its existing data might be stale. Data collection might be interrupted or not occur if the resource, network, or IBM Spectrum Control server is unavailable.

General Attributes	Descriptions of Triggering Conditions
Probe Status*	<p>One of the following statuses is detected for a probe:</p> <p>Not Successful</p> <p>An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning</p> <p>A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.</p> <p>Error (default)</p> <p>A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.</p> <p>For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Probe section on the Data Collection page.</p>
Status	<p>One of the following statuses is detected on a hypervisor:</p> <p>Not Normal</p> <p>An error or warning status was detected on the hypervisor or its internal resources.</p> <p>Warning</p> <p>A warning status was detected on the hypervisor or its internal resources.</p> <p>Error (default)</p> <p>An error status was detected on the hypervisor or its internal resources. For example, an error status might occur when a hypervisor goes offline, or a server disk is disconnected or partially disconnected if it has multiple paths and one of the paths is disconnected.</p> <p>Unreachable</p> <p>One or more of the monitored resources for a hypervisor are not responding. This status might be caused by a problem in the network.</p>

Controllers

Table 2. Triggering attributes and conditions for general changes on controllers

Controller Attributes	Descriptions of Triggering Conditions
New Disk Controller	A disk controller is detected for the first time. Use this alert to be notified of hardware additions on servers.
Removed Disk Controller	A previously monitored disk controller can no longer be found. Historical data about the controller is retained, but no current data is being collected. Use this alert to be notified if a controller is removed or becomes unavailable.

Data Stores

Table 3. Triggering attributes and conditions for general changes on data stores

Data Store Attributes	Descriptions of Triggering Conditions
Deleted Logical Volume	A previously monitored logical volume on a data store can no longer be found. Historical data about the volume is retained, but no current data is being collected. Use this alert to be notified if a volume is deleted or becomes unavailable.
New Logical Volume	A logical volume on the data store is detected for the first time.
VMDKs	You can use a number of operators to determine when you are notified of the number of virtual machine disk files (VMDKs) on the data store. For example, you can select to be notified when the number of VMDKs changes, or when the number is or is not a specific number, or when the number lies outside a specific range.

Table 4. Triggering attributes and conditions for capacity changes on data stores

Capacity Attributes	Descriptions of Triggering Conditions
Available Data Store Space	The amount of unused storage capacity on the hypervisor data store.
Available File System Capacity	The total amount of unused storage capacity on a file system.
Data Store Capacity	The total amount of storage capacity that is assigned to a data store.
File System Capacity (Previously known as Total File System Capacity)	The capacity of a file system.
Used File System Capacity	The used capacity on the hypervisor file system.
Used Capacity	The percentage of storage capacity that is used on the hypervisor data store.

Disks

Table 5. Triggering attributes and conditions for general changes on disks

General Attributes	Descriptions of Triggering Conditions
Firmware	The version of the Licensed Internal Code on the disk changes. You can use a number of operators to determine when you are notified of a firmware change, such as when the firmware is, or is not, a specific version, or when the version number contains a specific value.
Multipathing Policy	The multipathing policy that is in effect for a disk. For example, you can be notified when the policy changes, or when the policy is Round Robin, Load Balancing, Failover Only, or other another policy.
New Disk	A disk is detected for the first time. Use this alert to be notified of hardware changes on servers or hypervisors.
Paths	The number of access paths that are associated with the disk falls outside a specified range, or is equal to or not equal to a specified value.

General Attributes	Descriptions of Triggering Conditions
Removed Disk	A previously monitored disk can no longer be found. Historical data about the disk is retained, but no current data is being collected. Use this alert to be notified if a disk is removed or becomes unavailable.
Status*	One of the following statuses is detected on a disk: Not Normal An error or warning status was detected on the disk. Warning A warning status was detected on the disk. Error (default) An error status was detected on the disk.

Table 6. Triggering attributes and conditions for capacity changes on disks

Capacity Attributes	Descriptions of Triggering Conditions
Available Drive Capacity (Previously known as Available Disk Space)	The unused capacity on a hypervisor disk.
Capacity	The total amount of storage capacity assigned to a hypervisor disk.
Used Capacity	The amount of used storage capacity on a hypervisor disk .

Paths

Table 7. Triggering attributes and conditions for general changes on paths

Path Attributes	Descriptions of Triggering Conditions
Deleted Path	A previously monitored access path for a server disk can no longer be found. This change might or might not affect the availability of the disk because there might be more than one path available.
New Path	An access path for a disk is detected for the first time.
Status*	One of the following statuses is detected on a path: Not Normal An error or warning status was detected on the path. Warning A warning status was detected on the path. Error (default) An error status was detected on the path.

Virtual Machines

Table 8. Triggering attributes and conditions for general changes on virtual machines

Virtual Machine Attributes	Descriptions of Triggering Conditions
Deleted Virtual Machine	A previously monitored access path for a virtual machine can no longer be found. Historical data about the virtual machine is retained, but no current data is being collected. Use this alert to be notified if a virtual machine is deleted or becomes unavailable.
New Virtual Machine	A virtual machine is detected for the first time.
Status	One of the following statuses is detected on a virtual machine: Not Normal An error or warning status was detected on the virtual machine. Warning A warning status was detected on the virtual machine. Error (default) An error status was detected on the virtual machine.
VMDKs	You can use a number of operators to determine when you are notified of the number of VMDKs on the virtual machine. For example, you can select to be notified when the number of VMDKs changes, or when the number is or is not a specific number, or when the number lies outside a specific range.

Table 9. Triggering attributes and conditions for capacity changes on virtual machines

Capacity Attributes	Descriptions of Triggering Conditions
Capacity	The total amount of storage capacity assigned to a virtual machine hosted by the hypervisor.

VMDKs

Table 10. Triggering attributes and conditions for general changes on virtual machine disks (VMDKs)

VMDK Attributes	Descriptions of Triggering Conditions
Deleted Virtual Disk	A previously monitored virtual disk can no longer be found. Historical data about the virtual disk is retained, but no current data is being collected. Use this alert to be notified if a virtual disk is deleted or becomes unavailable.
New Virtual Disk	A virtual disk is detected for the first time.

Table 11. Triggering attributes and conditions for capacity changes on VMDKs

Capacity Attributes	Descriptions of Triggering Conditions
Size	The size of the VMDK on the hypervisor data store.

Capacity Attributes	Descriptions of Triggering Conditions
Used Capacity	The percentage of storage capacity that is used on the VMDK on the hypervisor data store. The used capacity is available if a Storage Resource agent is deployed.

Triggering conditions for switch alerts

You can set up IBM Spectrum® Control so that it examines the attributes and performance of a switch and notifies you when changes or violations are detected.

Alerts can notify you of general changes and performance issues on the following resources:

- [Switches \(performance\)](#)
- [Switches \(general changes\)](#)
- [Ports \(general changes\)](#)
- [Ports \(performance\)](#)
- [Trunks \(performance\)](#)

Important: Not all the attributes upon which you can alert are listed here. To view a complete list of attributes upon which you can alert, you can either edit alert definitions for a switch with no alert policy assigned, or you can create a new custom alert policy and define new alerts in the policy. To create a custom alert policy go to Configuration > Alert Policies and click Create Policy.

Not all the attributes upon which you can alert are listed here. A number of other attributes are available for alerts and are based on the key properties of a switch. To view a complete list of attributes upon which you can alert, go to Settings > Alert Policies. Double-click a default policy for a switch. Click Edit Alert Definitions on the Alert Definitions tab. View the attributes that are available in the general, capacity, and performance categories. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active.

In the tables, default alerts are marked with an asterisk (*).

Switches (general changes)

Table 1. Pre-defined alerts for Switches

General Attributes	Defining Conditions for Attributes
Last Successful Probe	A specified amount of time has passed since a probe or performance monitor was able to collect data about a switch. You can use this alert to be notified when up-to-date configuration, status, or performance data is not being collected about a switch and its existing data might be stale. This situation might occur if the resource, network, or IBM Spectrum Control server is unavailable.
Last Successful Monitor	
Probe Status*	<p>One of the following statuses is detected for a probe:</p> <p>Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.</p> <p>Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.</p> <p>For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Probe section on the Data Collection page.</p>
Performance Monitor Status*	<p>One of the following statuses is detected for a performance monitor:</p> <p>Not Normal An error or warning occurred during data collection. This status indicates that a performance monitor did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A performance monitor completed, but did not collect a complete set of performance data. This status might occur if the resource was rebooted during data collection, no valid performance data was provided by the resource, or a communication error occurred with the resource or its associated agent.</p> <p>Error A performance did not complete when it attempted to collect performance data about the resource. This status might occur if the resource cannot be reached during data collection, or if no configuration data is available for the resource.</p> <p>For details about why a specific status occurred, check the log for the performance monitor. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Performance Monitor section on the Data Collection page.</p>
Status*	<p>One of the following conditions is detected on a switch:</p> <p>Not Normal An error or warning status was detected on the switch or its internal resources.</p> <p>Warning A warning status was detected on the switch or its internal resources.</p> <p>Error (default) An error status was detected on the switch or its internal resources. For example, an error status might occur when a switch goes offline.</p> <p>Unreachable One or more of the monitored resources for a switch are not responding. This status might be caused by a problem in the network.</p>

Switches (performance)

Define alerts that notify you when the performance of a switch falls outside a specified threshold. In alerts, you can specify conditions based on metrics that measure the performance of switch ports, including I/O, data, and error rates, and frame transfer sizes. By creating alerts with performance conditions, you can be informed about potential bottlenecks in your network infrastructure.

For example, you can define an alert to be notified when the port congestion index for a port is greater than or equal to a specified threshold. Port congestion represents the estimated degree to which frame transmission was delayed due to a lack of buffer credits. Use this alert to help identify port conditions that might slow the performance of the resources to which those ports are connected.

You can also be notified when a metric is less than a specified threshold, such as when you want to identify ports that might be under used.

For a complete list of switch metrics that can be alerted upon, see [Performance metrics for switches](#).

Tips for performance conditions:

- A performance monitor must collect data about a resource before IBM Spectrum Control can determine whether a threshold is violated and an alert is generated for a performance condition.
- When you define a performance alert for the inter-switch connections category of a switch, the performance of ISL trunk connections is measured and alerted on. To alert on the performance of standard ISL connections, define performance alerts for the ports category of a switch.

Best practice: When you set thresholds for performance conditions, try to determine the best value so you can derive the maximum benefit without generating too many false alerts. Because suitable thresholds are highly dependent on the type of workload that is being run, hardware configuration, the number of physical disks, exact model numbers, and other factors, there are no easy or standard default rules.

A recommended approach is to monitor the performance of resources for a number of weeks and by using this historical data, determine reasonable threshold values for each performance condition. After that is done, you can fine-tune the condition settings to minimize the number of false alerts.

Click Edit alert Definitions, and for each performance alert definition click View History to see the history of switch, port or trunk performance and set the threshold you want relative to that data.

Trunks (performance)

Define alerts that notify you when the performance of a trunk falls outside a specified threshold. In alerts, you can specify conditions based on metrics that measure the performance of the trunk, including I/O, data, and error rates, and frame transfer sizes. By creating alerts with performance conditions, you can be informed about potential bottlenecks in your network infrastructure.

For example, you can define an alert to be notified when the aggregate port congestion index of the ports in the trunk is greater than or equal to a specified threshold. Port congestion represents the estimated degree to which frame transmission was delayed due to a lack of buffer credits. Use this alert to help identify port conditions that might slow the performance of the resources to which those ports are connected.

You can also be notified when a metric is less than a specified threshold, such as when you want to identify trunks that might be under used.

For a complete list of trunk metrics that can be alerted upon, see [Performance metrics for switches](#)

Ports (general changes)

Table 2. Pre-defined alerts for ports

General Attributes	Defining Conditions for Attributes
Removed Port	A previously monitored port can no longer be found. Historical data about the port is retained, but no current data is being collected. Use this alert to be notified if a port is removed or becomes unavailable.
State	A port is online, enabled but offline, or disabled.
Status*	One of the following statuses is detected for a port: Not Normal An error or warning status is detected on a port. Warning A warning status is detected on a port. This status might occur if the switch is stopped, starting, or in service (being maintained, cleaned, or administered). Error An error status is detected on a port.

Ports (performance)

Define alerts that notify you when the performance of a port falls outside a specified threshold. In alerts, you can specify conditions based on metrics that measure the performance of switch ports, including I/O, data, and error rates, and frame transfer sizes. By creating alerts with performance conditions, you can be informed about potential bottlenecks in your network infrastructure.

For example, you can define an alert to be notified when the port congestion index for a port is greater than or equal to a specified threshold. Port congestion represents the estimated degree to which frame transmission was delayed due to a lack of buffer credits. Use this alert to help identify port conditions that might slow the performance of the resources to which those ports are connected.

You can also be notified when a metric is less than a specified threshold, such as when you want to identify ports that might be under used.

For a complete list of port metrics that can be alerted upon, see [Performance metrics for switches](#)

Triggering conditions for fabric alerts

You can set up IBM Spectrum® Control so that it examines the attributes of a fabric and notifies you when changes are detected.

Alerts can notify you of general changes on fabric resources.

Important: Fabrics do not have a default alert policy assigned. If you wish to define alerts for a fabric you can either edit the alert definitions for a single fabric, or create an alert policy and assign it to one or more fabrics. To define alerts for a fabric, go to [Resources > Fabrics](#), right-click a fabric, and select [View Alert Definitions](#). Then, click [Edit Alert Definitions](#) to start defining alerts. To create an alert policy that can be applied to multiple fabrics, go to [Configuration > Alert Policies >](#) and click [Create Policy](#). Then, name the policy, click [Policy Type](#), select [Fabrics](#), and select the fabrics that you want to include.

Not all the attributes upon which you can alert are listed here. A number of other attributes are available for alerts and are based on the key properties of fabrics or switches. To view a complete list of attributes upon which you can alert, go to [Settings > Alert Policies](#). Double-click a default policy for a fabric or switch. Click [Edit Alert Definitions](#) on the [Alert Definitions](#) tab. View the attributes that are available in the general and capacity categories, and in the performance category for switches. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active.

Not all attributes are available for all types of fabrics. For example, you can't probe fabrics that contain Cisco switches or Brocade switches that run Fabric OS 8.2.1 or later. Therefore the [Last Successful Probe](#) and [Probe Status](#) attributes are not available for those types of fabric.

In the table, default alerts are marked with an asterisk (*).

Fabrics

Table 1. Alerts for fabrics

Fabric Attributes	Defining Conditions for Attributes
Probe Status*	<p>One of the following statuses is detected for a probe:</p> <p>Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.</p> <p>Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.</p> <p>For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Probe section on the Data Collection page.</p>
Status	<p>One of the following statuses is detected on a fabric:</p> <p>Not Normal An error or warning status was detected on the fabric or its internal resources.</p> <p>Warning A warning status was detected on the fabric or its internal resources.</p> <p>Error (default) An error status was detected on the fabric or its internal resources.</p> <p>Unreachable One or more of the monitored resources for a fabric are not responding.</p>
Custom alerts	You may choose from a small number of custom alerts to add to a fabric or policy. Conditions include number of Switches, number of Ports, etc.

Triggering conditions for server alerts

You can set up IBM Spectrum® Control so that it examines the attributes and capacity of a server and notifies you when changes are detected.

Alerts can notify you of general changes and capacity changes on the following resources:

- [Servers](#)
- [Controllers](#)
- [Disks](#)
- [Disk Groups](#)
- [File Systems and Logical Volumes](#)
- [Paths](#)
- [Shares](#)
- [Triggering conditions for the IBM Spectrum Control server](#)

Important: A number of other attributes are available for alerts and are based on the key properties of a server. To view a complete list of server attributes upon which you can alert, go to [Settings > Alert Policies](#). Double-click a default policy for a server. Click [Edit Alert Definitions](#) on the [Alert Definitions](#) tab. View the attributes that are available in the general and capacity categories. Note that the attributes that are automatically configured for alerts in the default alert policies, or default alerts, have a status of Active.

In the tables, default alerts are marked with an asterisk (*).

Tip: For capacity attributes, you can generate alerts when the amount of storage is greater than, less than, or equal to a specified value. You can also determine the unit of measurement for the attribute, such as KiB, MiB, GiB, or TiB.

Servers

Table 1. Triggering attributes and conditions for general changes on servers

General Attributes	Descriptions of Triggering Conditions
Agent State*	<p>A Storage Resource agent is in one of the following states:</p> <p>Not Normal An error or warning state was detected on a Storage Resource agent.</p> <p>Warning A warning state was detected on a Storage Resource agent. For example, a warning state might occur when an agent must be upgraded to the same version level as the IBM Spectrum Control server to which it is communicating.</p> <p>Error (default) An error state was detected on a Storage Resource agent. For example, an error state might occur when an agent was not able to be upgraded.</p>
Last Successful Probe	A specified amount of time has passed since a probe was able to collect data about a server. You can use this alert to be notified when up-to-date configuration and status data is not being collected about a server and its existing data might be stale. Data collection might be interrupted or not occur if the resource, network, or IBM Spectrum Control server are unavailable.
Probe Status*	<p>One of the following statuses is detected for a probe:</p> <p>Not Successful An error or warning occurred during data collection. This status indicates that a probe did not collect any data, or only collected a partial set of data about a resource.</p> <p>Warning A probe completed, but might not have collected a complete set of data. This status might occur if data cannot be collected about one or more of the internal resources of a resource.</p> <p>Error (default) A probe did not complete when it attempted to collect asset data about the resource. This status might occur if the resource cannot be reached during data collection.</p> <p>For details about why a specific status occurred, check the log for the probe. To check the log, go to the details page for a resource, click Data Collection, and select Actions > Open Logs in the Probe section on the Data Collection page.</p>
Status*	<p>One of the following statuses is detected on a server:</p> <p>Not Normal An error or warning status was detected on the server or its internal resources.</p> <p>Warning A warning status was detected on the server or its internal resources. For example, a warning status might occur when an HBA or HBA to a server node is newly discovered, is missing, or is rediscovered.</p> <p>Error (default) An error status was detected on the server or its internal resources. For example, an error status might occur when a server goes offline, or a server disk is disconnected or partially disconnected if it has multiple paths and one of the paths is disconnected.</p> <p>Unreachable One or more of the monitored resources for a server are not responding. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.</p>

Table 2. Triggering attributes and conditions for capacity changes on servers

Capacity Attributes	Descriptions of Triggering Conditions
Available Drive Capacity (Previously known as Available Disk Space)	The unused disk capacity on the local and SAN-attached storage for the server. SAN-attached storage is assigned to the server from storage systems.
Available File System Capacity	<p>The amount of unused capacity in the file systems on the server.</p> <p>Available file system capacity does not include capacity that is reserved for the operating system. For example, the available capacity for <code>tmpfs</code> on UNIX operating systems is not included in this value.</p>
Drive Capacity (Previously known as Total Disk Space)	The total disk capacity for all the local and SAN-attached storage on the server. SAN-attached storage is assigned to the server from storage systems.
File System Capacity	The amount of file system capacity on the server.
File System Capacity from Storage Systems	<p>The amount of file system capacity that is assigned to the server from storage systems.</p> <p>The file system capacity from storage systems is only available when SAN-attached storage is assigned to the server.</p>
Mapped SAN Capacity (Previously known as Assigned SAN Space)	<p>The amount of disk capacity that is assigned to the server from storage systems.</p> <p>The disk capacity from storage systems is only available when SAN-attached storage is assigned to the server.</p>
Used Capacity	The amount of used disk capacity on the local and SAN-attached storage for the server. SAN-attached storage is assigned to the server from storage systems.

Controllers

Table 3. Triggering attributes and conditions for changes on disk controllers

Controller Attributes	Descriptions of Triggering Conditions
<ul style="list-style-type: none"> Driver Version Firmware ROM Version 	<p>The version of the driver, firmware, or read-only memory (ROM) on a disk controller changes. You can use a number of operators to determine when you are notified of a version change, such as when the driver, firmware, or ROM is, or is not, a specific version, or when the version number contains a specific value.</p> <p>Use this alert for HBAs only.</p>
Last Data Collection	A specified amount of time since data was collected about a controller. Use this alert to be notified if data is not being collected about a controller or if the existing data is becoming too stale.
New Disk Controller	A disk controller is detected for the first time. Use this alert to be notified of hardware additions on servers.

Controller Attributes	Descriptions of Triggering Conditions
Removed Disk Controller	A previously monitored disk controller can no longer be found. Historical data about the controller is retained, but no current data is being collected. Use this alert to be notified if a controller is removed or becomes unavailable.
Status*	One of the following statuses is detected on a disk controller: Not Normal An error or warning status was detected on the controller. Warning A warning status was detected on the controller. Error (default) An error status was detected on the controller.

Disks

Table 4. Triggering attributes and conditions for general changes on disks

General Attributes	Descriptions of Triggering Conditions
Firmware	The version of the Licensed Internal Code on the disk changes. You can use a number of operators to determine when you are notified of a firmware change, such as when the firmware is, or is not, a specific version, or when the version number contains a specific value.
Multipathing Policy	The multipathing policy that is in effect for a disk. For example, you can be notified when the policy changes, or when the policy is Round Robin, Load Balancing, Failover Only, or other another policy.
New Disk	A disk is detected for the first time. Use this alert to be notified of hardware changes on servers or hypervisors.
Paths	The number of access paths that are associated with the disk falls outside a specified range, or is equal to or not equal to a specified value.
Removed Disk	A previously monitored disk can no longer be found. Historical data about the disk is retained, but no current data is being collected. Use this alert to be notified if a disk is removed or becomes unavailable.
Status*	One of the following statuses is detected on a disk: Not Normal An error or warning status was detected on the disk. Warning A warning status was detected on the disk. Error (default) An error status was detected on the disk.

Table 5. Triggering attributes and conditions for capacity changes on disks

Capacity Attributes	Descriptions of Triggering Conditions
Available Drive Capacity (Previously known as Available Disk Space)	The unused capacity on a disk that is attached to the server.
Capacity	The total amount of storage capacity on a disk that is attached to the server.
Used Capacity	The amount of used storage capacity on a disk that is attached to the server.

Disk Groups

Table 6. Triggering attributes and conditions for general changes on disk groups (volume groups)

General Attributes	Descriptions of Triggering Conditions
Deleted Volume Group	A previously monitored volume group can no longer be found. Historical data about the volume group is retained, but no current data is being collected. Use this alert to be notified if a volume group is deleted or becomes unavailable.
New Volume Group	A volume group is detected for the first time.
Status*	One of the following statuses is detected on a disk group: Not Normal An error or warning status was detected on the disk group. Warning A warning status was detected on the disk group. Error (default) An error status was detected on the disk group.

Table 7. Triggering attributes and conditions for capacity changes on disk groups (volume groups)

Capacity Attributes	Descriptions of Triggering Conditions
Available Capacity	The unused storage capacity on a server disk group.
Used Capacity	The amount of used storage capacity on a server disk group.
Volume Group Capacity	The total amount of storage capacity on a server volume group. This value is inclusive of all storage capacity and applies to all capacity values related to volume groups.

File Systems and Logical Volumes

Table 8. Triggering attributes and conditions for general changes on file systems and logical volumes

General Attributes	Descriptions of Triggering Conditions
--------------------	---------------------------------------

General Attributes	Descriptions of Triggering Conditions
Available Inodes	The number of unused inodes on file systems on the operating system changes, falls outside a specified range, or is equal to or not equal to a specified value.
Deleted File System	A previously monitored file system is deleted or unmounted from a server. Historical data about the file system is retained, but no current data is being collected. This attribute applies to file systems on the following resources: <ul style="list-style-type: none"> Storage systems that are configured for file storage, including Storwize® V7000 Unified Servers that are managed by Storage Resource agents
Deleted Logical Volume	A previously monitored logical volume can no longer be found. Historical data about the logical volume is retained, but no current data is being collected. Use this alert to be notified if a logical volume is removed or becomes unavailable.
New File System	A file system was detected for the first time. This alert applies to file systems on the following resources: <ul style="list-style-type: none"> Storage systems that are configured for file storage, including Storwize V7000 Unified Servers that are managed by Storage Resource agents
New Logical Volume	A logical volume is detected for the first time.
Used Inodes	The number of used inodes on file systems on the operating system changes. You can use a number of operators to determine when you are notified, such when the number of used inodes falls outside a specified range, or is equal to or not equal to a specified value.

Table 9. Triggering attributes and conditions for capacity changes on file systems and logical volumes

Capacity Attributes	Descriptions of Triggering Conditions
Available File System Capacity	The amount of unused storage capacity on a file system on the server disk.
File System Capacity	The total amount of storage capacity on a file system on the server disk.
Logical Volume Capacity	The total amount of storage capacity on a logical volume on the server disk.
Used File System Capacity	The capacity on a file system on the server disk.
Used Capacity	The percentage of used storage capacity on a file system or logical volume on the server disk.

Paths

Table 10. Triggering attributes and conditions for general changes on paths

Path Attributes	Descriptions of Triggering Conditions
Deleted Path	A previously monitored access path for a server disk can no longer be found. This change might or might not affect the availability of the disk because there might be more than one path available.
New Path	An access path for a disk is detected for the first time.
Status*	One of the following statuses is detected on a path: <ul style="list-style-type: none"> Not Normal <ul style="list-style-type: none"> An error or warning status was detected on the path. Warning <ul style="list-style-type: none"> A warning status was detected on the path. Error (default) <ul style="list-style-type: none"> An error status was detected on the path.

Shares

Table 11. Triggering attributes and conditions for general changes on shares

Share Attributes	Descriptions of Triggering Conditions
Deleted Share	A previously monitored share can no longer be found. Historical data about the share is retained, but no current data is being collected. Use this alert to be notified if a share is removed or becomes unavailable.
New Share	A new share was detected for the first time.

Triggering conditions for the IBM Spectrum Control server

The server on which IBM Spectrum Control is installed is automatically monitored for conditions that might cause an interruption in product functions. When these conditions are detected, alerts are triggered and shown on the Home > Alerts page. You do not need to manually define alerts for these product-related conditions; they are automatically enabled.

Table 12. Triggering conditions for the IBM Spectrum Control server

Triggering Condition	Explanation	Related Error Message
Database unavailable	The product database is not available. This database is the repository for information that is collected about the monitored resources in your environment.	ALR4112E, ALR4113E
High memory usage*	A high amount of memory is being used by a server process and might cause stability problems.	ALR4103W
Database alarm*	The system database or the database manager that hosts the product's database repository is reporting an alarm.	ALR4104W
High workload	The workload queue for the Device server is high and might cause performance issues.	ALR4105W
High number of external events	The server is receiving a high number of external events, such as CIM indications or SNMP traps. The high number of events might cause performance issues.	ALR4106W

Alert notifications and actions

Specify how you are notified when alert conditions are detected on resources, and define actions to take as a result of those alerts. These settings are defined globally for all resources, and can be overridden for a specific alert definition, for all alert definitions that apply to a specific resource, or for an alert policy.

Notifications and override behavior

You can specify notification settings at the following levels:

- Globally, for all alerts.
- In alert policies, for groups of resources of the same type.
- In individual alert definitions, which can be part of an alert policy, or can be specific to a resource.

Notification settings at the lower levels override the settings at the higher levels. For example, if you want to use email for alert notifications, you can specify email notification settings like this:

1. If you want all email notifications to go to a specific set of email addresses, specify email addresses for global alert notifications.
2. If you're using alert policies, you can also specify a list of email addresses to notify that are applicable to that policy. These addresses override the global setting.
3. For an alert definition, you can specify a set of email addresses to notify. This email addresses override the alert policy and global settings.

Scenarios for the different levels of alert notification

A small organization might specify only global email addresses if they want to send all alert notifications to one team.

A larger organization can have multiple storage environments that are managed by different teams. They can create alert policies for each environment and specify different email addresses for each alert policy. For example, they might create alert policies for their mainframe, VMware, and AIX® environments, and specify different email addresses for each of those policies. These email addresses override any global email addresses.

If you want to notify a specific person when a specific condition occurs, you set an email address in an individual alert definition. This email address overrides any alert policy or global email addresses. For example, an operations team might review all of the alerts, but might want to notify their storage architect if capacity in a pool exceeds 80%.

Specifying the frequency of alert notifications

You can use the Notification Frequency settings to avoid triggering too many alerts for some conditions. You can select one of the following options:

Option	Description
Send every time condition is violated	Receive alert notifications whenever an alert condition is violated.
Send once until problem clears	Receive one notification for a violation, even if the condition is violated multiple times.
Send every <i>time_period</i>	Receive one notification when an alert condition is initially violated. The alert is suppressed and no notifications are sent until both the specified time has passed and the alert is triggered again.

You can also select to only send notifications if the violation is not cleared for longer than a time period that you specify. Select Send after condition is violated for *time_period*, then specify the time period.

For example, you might want to be notified about an alert only after the condition has been violated for 20 minutes, and you only want to be notified about the alert once until the problem clears. In this case, you can set up an alert with the following notification frequency settings:

Notification Frequency

Send once until problem clears

▼

☒ Send after condition is violated for

20

↑
↓

minute(s)

▼

Restrictions:

- For some attributes, not all the notification frequency options are available. Specifically, you cannot change the notification frequency for the following attributes: New *resource*, Removed or Deleted *resource*, Last Successful Probe, and Last Successful Monitor.
- For attributes that use the changes operator (for example, the Firmware attribute), only the Send every time condition is violated and Send every *time_period* notification frequency options are available.

Specifying alert actions and notification settings in the GUI

Table 1. Locations where you can specify alert actions and notification settings

Task	Location in GUI
Specify the global notification settings for all alerts	Global Alert Notifications To specify the global notification settings, go to Configuration > Notification Settings, and specify the details for email, SNMP, and Tivoli® Netcool®/OMNIBus. These notification settings are applied globally to all alert definitions, unless overridden. Important: You must specify the global alert notifications settings for each type of alert notification that you want to enable.

Task	Location in GUI
Specify the notification settings for an alert policy	Alert policies page To specify the notification settings for an alert policy, go to Settings > Alert Policies. Double-click the policy whose notification settings you want to specify. Click Edit Policy Notifications. Specify the details for email, SNMP, and Tivoli Netcool/OMNIBus. These notification settings are applied to all of the alert definitions for all resources in the policy, unless overridden.
Specify the notification settings for all the alert definitions that apply to a resource that is not managed by an alert policy	Resource details page For example, to specify the notification settings for a block storage system, go to Storage > Block Storage Systems. Double-click the resource for which you want to define alerts. In the General section of the resource details page, click Alerts. Then click Edit Notifications. These notification settings are applied to all the alert definitions that are specified for the selected storage system. You can override the global settings by specifying different settings for each alert definition.
Specify alert actions and notification settings for a specific alert definition	Alert definitions page, Resource details page For example, to specify the notification settings for an alert definition in an alert policy, go to Settings > Alert Policies. Double-click the policy whose notification settings you want to specify. Click Edit Alert Definitions. Alternatively, to specify the notification settings for an alert definition for a resource that is not managed by an alert policy, go to Storage > Block Storage Systems. Double-click the resource for which you want to define alerts. In the General section of the resource details page, click Alerts. Then click Edit Alert Definitions. Click on the alert definition to show advanced options. Specify the details for email, SNMP, and Tivoli Netcool/OMNIBus. The notification settings that you define override any global notification settings, any policy settings, and any settings for the resource.

Alert actions

You can specify that the following actions are taken when alert conditions are detected on monitored resources:

Run script

Run a script when an alert is triggered for the condition. Use a script to call external programs or run commands that take action as the result of an alert. By using a script, you can automatically address potential storage issues when they are detected to avoid unplanned downtime or performance bottlenecks. [Learn more.](#)

Email

Send emails to specific email addresses when an alert is detected on a monitored resource.

Netcool or OMNIBus EIF Probe Server

Send alert notifications to a Netcool server or OMNIBus EIF probe server within your environment that was configured to receive IBM Spectrum Control alerts.

SNMP trap

Generate SNMP trap messages to any network management station (NMS), console, or terminal when an alert condition is detected. System administrators must set up their SNMP trap ringer with the provided management information base (MIB) files to receive SNMP traps from the product.

Windows event log or UNIX syslog

Write alert messages to the OS log. If you already have an administrator monitoring OS logs, this method is a way to centralize your priority messages for quick notification and viewing.

- [How scripts are run](#)

After an alert is triggered, scripts can be run from the server where IBM Spectrum Control is installed or from any server where a Storage Resource agent is deployed. For some alert conditions, scripts can be specified to run on the server where the change was detected.

Related tasks

- [Defining alerts](#)

How scripts are run

After an alert is triggered, scripts can be run from the server where IBM Spectrum® Control is installed or from any server where a Storage Resource agent is deployed. For some alert conditions, scripts can be specified to run on the server where the change was detected.

You must include the script files that you want to run in the /scripts directory on the server where IBM Spectrum Control is installed. The default path of /scripts is determined by the operating system of the server:

- Windows: C:\Program Files\IBM\TPC\Data\scripts
- UNIX or Linux®: /opt/IBM/TPC/Data/scripts or /usr/IBM/TPC/Data/scripts

Tip: When you run a script against a NAS filer, the script is run from the Storage Resource agent that is assigned to the file system where the condition was detected.

Script parameters

You can run scripts in response to alerts. Parameters for the conditions and attributes in an alert definition are included in the script. Each parameter is associated with a number, such as \$1, \$2, \$3, and so on. These parameters are used when the alert is triggered.

The following examples show the parameters that are included in a script for different alerting scenarios.

Alerting example: Storage System Pool Used Capacity >= 10

```
$1 = Pool
$2 = tpcvm4-3
$3 = "Storage System"
$4 = SVC-storage1
```



```
$5 = "Storage System Pool Used Capacity Percent >= 10 %"
$6 = "75.97 %"
```

Alerting example: Storage System Write Data Rate <= 1,100 MiB/s

```
$1 = ""
$2 = ""
$3 = "Storage System"
$4 = SVC-storage1
$5 = "Write Data Rate <= 1,100 MiB/s"
$6 = "0 MiB/s"
```

Alerting example: Pool Total I/O Rate <= 445 ops/s

```
$1 = Pool
$2 = scv5k2c0_pool
$3 = "Storage System"
$4 = SVC-storage1
$5 = "Total I/O Rate - overall <= 445 ops/s"
```

Alerting example: Custom alert with the multiple conditions

Conditions:

- Storage System Volume Capacity >= 1,000,000 KiB
- Storage System Volume Compressed is yes
- Storage Volume Compression Savings Percent >= 13

Parameters:

```
$1 = Volume
$2 = vol1
$3 = "Storage System"
$4 = SVC-storage1
$5 = "Storage System Volume Capacity >= 1,000,000 KiB"
$6 = "2,097,152 KiB"
$7 = Volume
$8 = vol1
$9 = "Storage System"
$10 = SVC-storage1
$11 = "Storage System Volume Compressed is yes"
$12 = yes
$13 = Volume
$14 = vol1
$15 = "Storage System"
$16 = SVC-storage1
$17 = "Storage Volume Compression Savings Percent >= 13 %"
$18 = "17.25 %"
```

Restriction: For some alerts, such as custom alerts with multiple conditions and attributes, the number of parameters in the associated script might exceed nine. In Windows, only the first nine parameters in a script can be referenced by number (parameter %1 to parameter %9). Additional parameters are ignored.

How script names are resolved

When you run scripts in response to alerts, the names of those scripts are resolved by using the following criteria:

- The operating system of the server where the product is installed.
- The script name that you specify on the Script name field when you define or edit an alert definition.

The operating system of the server or Storage Resource agent where you choose to run a script determines how that script that is run:

UNIX, Linux operating systems

A Storage Resource agent that is deployed on a UNIX or Linux operating system does not run scripts that have an extension. If the specified script name contains an extension, the agent ignores the extension and searches for a file of the same name (without an extension). For example, if you enter **backup.vbs** in the Script name field, a UNIX or Linux agent searches for a file with the name **backup**. If the script exists, the first line of the script is read and the appropriate interpreter is used to run the script. If a file cannot be found, no action is taken.

Windows operating systems

A Storage Resource agent that is deployed on a Windows operating system runs scripts that have an extension. The extension of a script file determines which interpreter is used to run the script. If you specify a script name that does not contain an extension, the agent searches for the file name in the following order (based on its extension): **.com**, **.exe**, **.bat**, **.cmd**, **.vbs**, **.vbe**, **.js**, **.jse**. If a file cannot be found, no action is taken.

Tip: You can run scripts with the same name across multiple operating systems. To run a script on both a UNIX or Linux and Windows operating systems, you must have two versions of that file in the /scripts directory. For example, if you want to run a provisioning script against Storage Resource agents on UNIX or Linux and Windows servers, you must have two versions of that file (**provision.bat** and **provision**) in /scripts.

Viewing information about resources

You can view detailed information about the resources that are monitored by IBM Spectrum® Control. Resources include storage systems, volumes, pools, servers, hypervisors, fabrics, and switches. You can also view information about internal and related resources.

Remember: Data collection jobs retrieve information about monitored resources. To ensure that the most up-to-date information about resources is available in the GUI, run data collection jobs regularly. To learn about how to collect data, see [Collecting data](#).

- [How information is organized](#)

Information about monitored resources is organized into different pages in the GUI. You can use the pages to view summary information about all types of

monitored resources, such as storage systems, fabrics, and switches. You can see what resources are consuming storage and the alerts that are being detected on those resources. You can also view a chart of the most active storage systems in your environment.

- [Dashboard view](#)

Use the dashboard to monitor the overall condition of monitored resources and identify potential problem areas in a storage environment. The dashboard also displays alerts, as well as a chart showing the most active storage systems.

- [Viewing information about top-level resources](#)

Use a resource list page as a central location for viewing information about all the monitored resources of a specific type. Resources that have a list page include storage systems, volumes, pools, shares, servers, hypervisors, switches, and fabrics.

- [Viewing information about internal, object, and related resources](#)

Use resource details pages to view information about a specific top-level resource and its internal, object, and related resources.

- [Overview charts](#)

Monitor the activity and storage usage of your storage systems. Use the charts and key metrics to gauge the workload activity of your storage resources and to check whether you have sufficient capacity.

- [Viewing information about enclosures](#)

View information about enclosures. Enclosures are the metal structures in which various storage components are mounted, including control units, nodes, disks and drives, and power supplies.

How information is organized

Information about monitored resources is organized into different pages in the GUI. You can use the pages to view summary information about all types of monitored resources, such as storage systems, fabrics, and switches. You can see what resources are consuming storage and the alerts that are being detected on those resources. You can also view a chart of the most active storage systems in your environment.

Dashboard view

The dashboard is view displayed when you start a new session of IBM Spectrum® Control or when you go to Home > Dashboard in the menu bar. Use the dashboard to view the following summary information about a storage environment

- Condition of storage: block, file, and object storage systems
- Condition of storage consumers: hypervisors and servers
- Condition of network resources: fabrics and switches
- Number error, warning, and information alerts that are generated on resources
- Most active storage systems in your environment.

You can click on resource icons in the dashboard to add more resources for monitoring and to see detailed information about resources that IBM Spectrum Control is currently monitoring. Position the mouse pointer over the lines in the Most Active Storage Systems chart to see performance information about specific resources. For more information about the dashboard, see [Dashboard view](#).

Resource list pages

Use a resource list page as a central location for viewing information about the top-level resources of a specific type. For example, use the File Storage Systems page to view a list of the file storage systems that are monitored by IBM Spectrum Control. Each resource list page corresponds to a specific type of resource, which includes block storage systems, file storage systems, object storage systems, volumes, pools, servers, hypervisors, switches, and fabrics. You can view the following information on resource list pages:

- Overall condition of the resources
- Key information about the resources
- Alerts that were generated for the resources
- Jobs that were run for the resources

To learn about how to access the list pages for resources, see [Viewing information about top-level resources](#).

Note: Volumes and pools are internal resources of storage systems, but have their own resource list pages so that you can view them all in a central location.

Resource detail pages

Use resource details pages to view information about resources that are internal or related to a top-level resource. For example, use the File Storage System details page to view detailed information about a specific file storage system. You can view the following information on resource details pages:

- The overall capacity and storage usage for a resource
- The status of resources
- Key information about a resource
- Alerts that were generated for a resource
- Data collection that was run for a resource
- Information about components that exist in top-level resources
- Information about the resources that are related to another resource

To learn about how to access detail pages for resources, see [Viewing information about internal, object, and related resources](#).

Alerts page

Use the Alerts page to view and manage the alert conditions that were detected on monitored resources. Specifically, you can complete the following actions:

- View the overall number of alerts that were detected on monitored resources. This information is presented in a bar chart for the following types of alerts:
 - Database
 - Fabric
 - Hypervisor

- Other
- Performance
- Server
- Storage system
- IBM Spectrum Control
- View information about specific alerts, including severity, last occurrence, and the resources where they were detected.
- Acknowledge alerts that were reviewed but are not yet resolved.
- Unacknowledge alerts.
- Remove alerts from the page.

To access the Alerts page, go to Home > Alerts in the menu bar.

Performance Monitors page

IBM Spectrum Control can collect information about the performance of monitored storage systems and switches. This information includes key performance metrics. On the Performance Monitors page, you can view detailed information about performance monitors that are collecting data. Use this information to identify performance monitors that are encountering problems when they run. You can also view the metrics that are collected by a performance monitor for a resource.

To access the Performance Monitors page, select Home > Performance Monitors.

Tasks page

Use the Tasks page to view and manage the tasks that IBM Spectrum Control uses to optimize resources and provision storage. Tasks are created when you complete any of the following actions:

- Analyzing tiering by moving volumes to higher or lower tiers.
- Transforming storage by moving or converting volumes.
- Balancing pools by distributing the workload of volumes across pools on the same tier.

To access the Tasks page, go to Home > Tasks .

System Management page

Use the System Management page to view information about the overall condition of IBM Spectrum Control. You can view information about the servers on which IBM Spectrum Control is installed. This information includes component server and database status, certain server alerts and database connection alerts, server file-system capacity information, and remote volume-performance information.

To access the System Management page, go to Home > System Management in the menu bar.

Reporting

Use the optional Cognos® Analytics reporting tool in the GUI to view detailed reports about resources. The following reports are available:

Predefined reports

These reports are included with IBM Spectrum Control. The following types of predefined reports are available:

- Capacity and relationships of resources
- Performance
- Historical space
- Storage tiering

Custom reports

Unlike predefined reports, you can select the information that you want to include in custom reports. Custom reports can contain detailed information about the relationships between monitored resources, the properties of monitored resources, and detailed information about the performance of monitored resources.

To learn about how to use the Cognos Analytics reporting tool, see [Reporting with Cognos Analytics](#).

Dashboard view

Use the dashboard to monitor the overall condition of monitored resources and identify potential problem areas in a storage environment. The dashboard also displays alerts, as well as a chart showing the most active storage systems.

Viewing the status of monitored resources

The dashboard provides status information for the resources that IBM Spectrum® Control is monitoring.

The overall status for each resource type

A status symbol and number are shown below each resource icon on the dashboard. This symbol represents the most critical status that was detected on the internal resources for a resource type. For example, if 20 storage systems are monitored, and an error was detected on a volume for one of those storage systems, a



1 status is shown below the storage systems icon or label in the dashboard. If no errors, warnings, or unreachable statuses were detected on the monitored resources, then a green symbol is shown without a number.

Use the status to quickly determine the condition of your monitored resources and if any actions must be taken. On the dashboard, the following statuses might be shown for a resource type:



Error (red)

An error status was detected on one or more of the monitored resources for a resource type. Error statuses represent serious problems on a resource. Resolve these problems as soon as possible. View the details of a resource to learn more about the status of its internal resources.



Unreachable (orange)

One or more of the monitored resources for a resource type are not responding. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.



Warning (yellow)

A warning status was detected on one or more of the monitored resources for a resource type. Warning conditions are not critical, but represent potential problems. View the details of a resource to learn more about the status of its internal resources.



Normal (green)

No warnings or errors were detected on the monitored resources for a resource type.



Unknown status (gray)

One or more of the monitored resources for a resource type have an unknown status. This status might occur if no data was collected about a resource. To change an unknown status, use a data collection job to collect status information about the resource.

For information about how status is determined, see [How the condition of a resource is determined](#).

The number of statuses that were detected and acknowledged for each resource type

To view the total number of statuses for a resource type, move the mouse pointer over the related status symbol. Hover help shows the total number of occurrences for each status, and the number of statuses that were acknowledged by a IBM Spectrum Control user. An acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.

Viewing the status of monitored storage systems

Storage systems can be configured for block storage, file storage, object storage, a combination of block and file storage, or a combination of file and object storage. The top left section of the dashboard shows the type, number, and status of the storage systems that IBM Spectrum Control is monitoring. The number in the circle next to the storage system type shows how many storage systems of that type are being monitored. The status icon below each storage system type shows the status of the monitored systems. Click the storage system type, for example Block Storage Systems, to see more detailed status and information about those storage systems. Use the information in the horizontal bar charts to view details about the used and free space of the monitored storage systems in your environment:

Storage systems that are configured for block data

The following information is shown for monitored storage systems that are configured for block data:

Used Block Storage Space

Amount of space that is being used for block data on all the storage systems that are configured for block data.

Available Block Storage Space

Amount of unused space on all the storage systems that are configured for block data.

Storage systems that are configured for file data

The following information is shown for monitored storage systems that are configured for file data:

Used File System Capacity

Amount of file system capacity that is being used by files and directories on the storage systems and filers that are configured for file data.

Available File System Capacity

Amount of unused storage space on the storage systems and filers that are configured for file data.

Storage systems that are configured for object data

The following information is shown for monitored storage systems that are configured for object data:

Used Object Storage Space

Amount of storage space that is used on all the object storage systems.

Available Object Storage Space

Amount of storage space that is available on all the object storage systems.

Viewing the status of monitored resources that are consuming storage

The diagram in the top middle section of the dashboard provides a visual representation of the monitored resources that are consuming storage. You can click the icons in this section to add resources of that type or to see detailed information about the resources you have already added. The number next to each icon shows the number of those resources that are being monitored by IBM Spectrum Control. If IBM Spectrum Control is monitoring resources of a particular type, such as servers, you will see a status icon that shows the overall status of those types of resources.

The types of resources that are consuming storage

Each type of resource in the diagram is represented by an icon. To learn more detailed information about the individual resources that are monitored, click the icon in the diagram. For example, if you click the departments icon, a new page shows detailed information about the individual departments that are being monitored.

The number of resources that are monitored for each resource type

A number in a circle is shown next to each resource icon. This number represents the number of monitored resources for each resource type. For example, if 4 file storage systems are monitored, the value ④ is shown. This number excludes resources that are not monitored by IBM Spectrum Control.

Tip: If no resources are being monitored for a resource type, click its icon on the dashboard to add a resource of that type. For example, if the dashboard shows a 0 next to the departments icon, you can click the icon or the circle for departments to add a department. The GUI guides you through the process for adding the resource for monitoring. For more information about adding resources, see [Adding resources](#).

Viewing the status of monitored network fabrics and switches

The top right section of the dashboard shows the number and status of the fabrics and switches that IBM Spectrum Control is monitoring. The number in the circle next to the fabric or switch icon shows how many fabrics and switches are being monitored. The status icon below the fabric or switch shows the status of those entities. Click the icon to see more detailed status and information about those fabrics and switches. The dashboard shows the following information for the monitored fabrics and switches in your environment:

Fabrics

- The number of fabrics that are being monitored.
- The most critical status of the monitored fabrics. If more than one fabric is monitored, the dashboard shows the most critical status level of those fabrics.

Switches

- The number of switches that are being monitored.
- The most critical status of the monitored switches. If more than one fabric is monitored, the dashboard shows the most critical status level of those switches.

Viewing the status of alerts that were detected on monitored resources

Alerts are generated when IBM Spectrum Control detects certain conditions or events on monitored resources. The Alerts section of the dashboard shows a status summary of the alerts that were detected in your environment during a specified time period. This summary includes the number of alerts that have the following statuses: Critical, Warning, and Informational.

Use the following actions to manage the display of statuses for alerts:

- To select the time period in which you want to view alert statuses, click the down arrow and select a time period. You can view the status of all alerts that were detected in the last hour, last day, or last week.
- To view more detailed information about alerts, click View all alerts to access the Alerts page.

Viewing the most active storage systems in your environment

The Performance chart provides information about the most active storage systems in your storage environment. Position your cursor over a data point in the chart to see the total I/O rate for a particular storage system. The I/O rate is displayed in operations per second (ops/s) on the left vertical axis, and in milliseconds per operation (ms/ops) on the right vertical axis. The horizontal axis shows the time of the performance data in hourly increments.

Viewing information about top-level resources

Use a resource list page as a central location for viewing information about all the monitored resources of a specific type. Resources that have a list page include storage systems, volumes, pools, shares, servers, hypervisors, switches, and fabrics.

About this task


Procedure


1. In the menu bar, hover the mouse pointer over the menu item for the type of resource that you want to view.
You can access the following types of resources:
 - Storage Block storage systems, file storage systems, Object Storage systems, volumes, pools, shares
 - Servers Servers, hypervisors
 - Network Switches, fabrics
2. Select the type of resource that you want to view.
For example, if you want to view information about servers, go to Servers > Servers.
A list of monitored servers is displayed on Servers page.
3. Locate the resource that you want to view in the list.
Each column in the list shows information about the resource.
4. Optional: To view key attributes of a specific top-level resource without leaving the page, right-click the resource and select View Properties.
A properties notebook is displayed. Information in a properties notebook can include storage statistics, hardware attributes, volume assignments, disk assignment, and performance metrics.
5. Optional: To view more detailed information about a specific top-level resource, including its internal and related resources, right-click the resource and select View Details.
A resource details page is displayed. The page includes information about storage statistics, alerts, jobs, internal resources, and related resources.

Results

A resource list page shows information about the monitored resources for the selected type. You can view the following information on the page:

- The overall status of the resources for a specific type.
- Key information about the resources.
- Alerts that were generated for the resources.
- Jobs that were run for the resources.

Tip: To view descriptions of the information that is available on a resource list page, click the Help icon  in the upper right of the page.

To view information about rollup resources, you must be in rollup mode. To enter rollup mode, move the mouse pointer over the rollup icon  in the menu bar and select Enter rollup mode.

Related reference

- [Resources that you can monitor](#)

Viewing information about internal, object, and related resources

Use resource details pages to view information about a specific top-level resource and its internal, object, and related resources.

About this task


Procedure

1. In the menu bar, hover the mouse pointer over the menu item for the type of resource that you want to view.
You can access the following types of resources:
 - Storage: block storage systems, file storage systems, object storage systems, volumes, pools, shares
 - Servers: servers, hypervisors
 - Network: switches, fabrics
2. Select the type of the resource that you want to view.
A resource list page shows the monitored resources for the selected type. For example, if you want to view a block storage system, go to [Storage > Block Storage Systems](#). A list of monitored block storage systems is displayed.
3. In the list of resources, right-click the resource that you want to view and select View Details.
A resource details page is displayed. This page includes information about storage statistics, alerts, jobs, internal resources, object resources, and related resources.
4. To view information about the internal resources of the top-level resource, click the name of the internal resource in the Internal Resources section.
A list of the internal resources is displayed. Each column in the list shows information about the resources. For example, on the details page for a server, click Controllers to view a list of the controllers that are associated with the server.
5. Optional: To view key attributes of a specific internal resource without leaving the page, right-click the resource and select View Properties.
6. To view information about the object resources in an object storage system, click the name of the object resource in the Object Resources section.
Object resources are associated only with object storage systems, such as IBM Spectrum Scale that is configured for object storage.
7. To view information about resources that are related to the top-level resource, click the name of the related resource in the Related Resources section.
A list of the related resources is displayed. For example, on the server details page, click Block Storage Systems to view a list of the block storage systems that have at least one disk (volume) assigned to the server.
8. Optional: To view key attributes of a specific related resource without leaving the page, right-click the resource and select View Properties.
9. Optional: To view the resource details page for a related resource, right-click the resource and select View Details.

Results

A resource details page shows information about a top-level resource. The information is available in the following links and sections on the page:

- Overview: The overall capacity and storage usage for a resource.
- Properties: Key information about a resource.
- Alerts: Alerts that were generated for a resource.
- Tasks: Jobs that were run for a resource.
- Internal Resources: Information about the components that exist in a higher-level resource.
- Object Resources: Information about the resources in a IBM Spectrum Scale cluster, such as accounts and containers, that enable the IBM Spectrum Scale cluster to be used as an object storage system.
- Related Resources: Information about the resources that are related to another resource. A related resource is external to a higher-level resource, but is related to it through assigned storage, a network connection, or virtual hosting.

Tip: To view descriptions of the information that is available on a resource details page, click the Help icon  in the upper-right corner of the page.

Related reference

- [Resources that you can monitor](#)

Overview charts

Monitor the activity and storage usage of your storage systems. Use the charts and key metrics to gauge the workload activity of your storage resources and to check whether you have sufficient capacity.

Review the charts and key information about the activity and availability of capacity:

- [For block storage systems](#)
- [For file storage systems](#)
- [For object storage systems](#)

Block storage systems

On the Overview page for a storage system, monitor storage usage, detect capacity shortages, and plan your capacity needs. The type of block storage system determine the information that is displayed.

Metrics for storage systems that have both block and file storage

The following information is shown for storage systems that have both block and file storage, such as Dell EMC Unity and NetApp ONTAP 9:


- The Capacity section shows how much capacity is used and how much capacity is available for storing data.
- The Provisioned Capacity section shows the written capacity values in relation to the total provisioned capacity values before data reduction techniques are applied:
 - % Written: The capacity of the data that is written to the volumes as a percentage of the total provisioned capacity of the volumes.
 - GiB Available: The amount of capacity that is still available for writing data to the thin-provisioned volumes. It is the difference between the provisioned capacity and the written capacity, which is the thin-provisioning savings.

- The Capacity Savings section shows the amount of capacity that is saved by using deduplication, thin provisioning, and the data compression capabilities of the storage system.

Get more views of Capacity, Provisioned Capacity, and Capacity Savings:

- To get a breakdown of the capacity usage by pool or volume, click the links in each section. For example, on the Block Storage Systems page, double-click a storage system, and then click View capacity by pool. You get a capacity breakdown of the storage system by pool, and you see the fill and growth rate for each pool.
- To view the capacity savings of a storage system with data reduction metrics and savings summary that includes duplication, thin provisioning, drive compression, and pool compression, click View capacity savings. If a capacity limit was set, you can hover over the capacity chart and check how much available capacity you have left before you reach the capacity limit.

Figure 1. Capacity and activity charts



Three types of charts are provided:

- Capacity overview charts
- Capacity by storage resource charts
- Workload activity charts

The Capacity chart at the top of the Overview page shows how much capacity is used and how much capacity is available for storing data.

The Provisioned Capacity chart shows the written capacity values in relation to the total provisioned capacity values before data reduction techniques are applied. The following values are shown:

- The capacity of the data that is written to the volumes as a percentage of the total provisioned capacity of the volumes.
- The amount of capacity that is still available for writing data to the thin-provisioned volumes in relation to the total provisioned capacity of the volumes. Available capacity is the difference between the provisioned capacity and the written capacity, which is the thin-provisioning savings.

A breakdown of the total capacity savings that are achieved when the written capacity is stored on the thin-provisioned volumes is also provided.

In the capacity overview chart, a horizontal bar is shown when a capacity limit is set for the storage system. Just hover over the chart to find out what the capacity limit is and how much capacity is left before the capacity limit is reached.

To get a breakdown of the capacity usage by pool or volume, click the links.

Try it: Click View capacity by pool. You get a capacity breakdown of the storage system by pool, and you can see the recent fill and growth rates for each pool in the storage system.

Depending on the type of storage system, the following activity and capacity by storage resource charts are shown:

Table 1. Capacity by storage resource

Chart Name	Purpose
Capacity	Use this chart to monitor the available, used, and provisioned capacity in your storage system over the last 30 days.
Capacity by Host	Use this chart to monitor the current capacity of the storage system that is mapped to hosts.
Capacity by Pool	Use this chart to monitor the current available and used capacity of your largest pools.
Capacity by Volume	Use this chart to monitor the current used and unused provisioned capacity of your largest volumes.
Capacity by Tier	Use this chart to monitor the current used and available capacity of the pools that are assigned to tiers.

Exclusive to IBM Storage Insights Pro: In IBM® Storage Insights Pro, users with Administrator privileges, can set capacity limits for storage systems and pools. When you hover over the capacity overview chart, you can see the value for the Capacity Limit and the Capacity-to-Limit. The Capacity-to-Limit is the amount of capacity that is left before the capacity limit is reached.

Table 2. Activity charts

Chart Name	Purpose
Overall system activity	Use this chart to monitor the overall activity of your storage system, which is broke down into total I/O rate, read I/O rate, and write I/O rate in operations per second over the last 24 hours.
Most Active Nodes	Use this chart to monitor the nodes with the heaviest workloads, which is calculated in operations per second, over the last 24 hours.
Most Active Pools	Use this chart to monitor the pools with the heaviest workloads, which is calculated in operations per second, over the last 24 hours.
Most Active Volumes	Use this chart to monitor the volumes with the highest response times over the last 24 hours.
MDisk Activity	Use this chart to monitor managed disks (M Disks) with the heaviest workloads, which is calculated in operations per second, over the last 24 hours.

File storage systems

The following capacity charts and Inode usage charts are shown for file storage systems:

Table 3. Capacity and Inode charts

Chart Name	Purpose
Total File System Space	Use this chart to monitor the capacity, the used capacity, and the available capacity of your file systems over the last 30 days.
Capacity by File System	Use this chart to monitor the current capacity and the used and available capacity of your fullest file systems .
Inodes by File System	Use this chart to monitor the availability of Inodes.

Object storage systems

The following capacity charts are shown for object storage systems:

Table 4. Charts for object storage systems

Chart Name	Purpose
Capacity	Use this chart to monitor the total capacity, used and available capacity of your object storage system.
Sites	Use this chart to monitor the distribution of Accesser® and SliceStor® nodes across sites.
Capacity by IDA	Use this chart to monitor the used capacity of the vaults that are configured with IDA (Information Dispersal Algorithm).
Failure Tolerance	Monitor the availability and fault tolerance of vaults and SliceStors.

- [Capacity by pool](#)
Monitor the capacity, fill rate %, and growth rate for your pools. You can also check how much capacity is left before the pools reach their capacity limit.
- [Capacity by volume](#)
Monitor the total provisioned capacity of your volumes to identify the volumes that have sufficient capacity to meet your storage needs.

Capacity by pool

Monitor the capacity, fill rate %, and growth rate for your pools. You can also check how much capacity is left before the pools reach their capacity limit.

To help you make decisions about the capacity that you need to store data, monitor the growth in used capacity and the following key capacity metrics:


Adjusted Used Capacity (%)

The amount of capacity that can be used without exceeding the capacity limit.

Example: Adjusted Used Capacity

Before Capacity Limit Was Set

100 GiB Capacity



■ Used Capacity = 40 GiB

■ Available Capacity = 60 GiB

After Capacity Limit Was Set

100 GiB Capacity



■ Capacity Limit = 80 GiB or 80 GiB

■ Adjusted Used Capacity = 50 GiB or 40 GiB

■ Capacity-to-Limit = 30 GiB or 40 GiB

The formula for calculating Adjusted Used Capacity (%) is $(\text{Used Capacity in GiB} / \text{Capacity Limit in GiB}) * 100$. For example, if the capacity is 100 GiB, the used capacity is 40 GiB, and the capacity limit is 80 GiB, then the value for Adjusted Used Capacity (%) is $(40 \text{ GiB} / 80 \text{ GiB}) * 100$ or 50%. So, in this example, you can use 30% or 40 GiB of the usable capacity of the resource before you reach the capacity limit.

If the used capacity exceeds the capacity limit, the value for Adjusted Used Capacity (%) is over 100%.

To add the Adjusted Used Capacity (%) column, right-click any column heading on the Pools page.

See these related values for more information Capacity Limit (%) and Capacity-to-Limit (GiB).

Availability: This metric is not available for all storage systems, such as Dell EMC VMAX.

Available Capacity (GiB)

The amount of usable capacity that is not yet used in the pool.

Capacity (GiB)

The amount of capacity that is available for storing data in the pool after formatting and RAID techniques are applied.

Capacity Limit (%) and Capacity Limit (GiB)

The limit that was set on the capacity that is used by your pools. For example, the policy of your company is to keep 20% of the usable capacity of your pools in reserve. So, you log into the GUI as Administrator and set the capacity limit of your pools to 80%.

Example: Administrator Sets Capacity Limit to 80%

100 GiB Capacity



■ Used Capacity = 40 GiB

■ Available Capacity = 60 GiB

100 GiB Capacity



■ Capacity Limit = 80 GiB or 80 GiB



Click the illustration above to find out how to set capacity limits.

The GiB value for the capacity limit for the pool is calculated when you set the value for the Capacity Limit (%).

To add the Capacity Limit (%) and the Capacity Limit (GiB) columns, right-click any column heading on the Pools page.

See these related values for more information Adjusted Used Capacity (%) and Capacity-to-Limit (GiB).


Zero capacity: When you set the capacity limit for pools, the values shown for Zero Capacity are readjusted to take into account the capacity limit of the pool. The date will represent when the capacity limit of the pool is reached. If the pool has already reached the capacity limit, Depleted is shown. None is shown when a trend in storage consumption can't be detected because the pool's storage isn't being consumed or because not enough data was collected to predict storage consumption.

Availability: This metric is not available for all storage systems, such as Dell EMC VMAX.

Capacity-to-Limit (GiB)

The amount of capacity that is available before the capacity limit is reached.

Example: Capacity-to-Limit



The formula for calculating Capacity-to-Limit (GiB) is (Capacity Limit in GiB - Used Capacity in GiB). For example, if the capacity limit is 80% or 80 GiB and the used capacity is 40 GiB, then the value for Capacity-to-Limit (GiB) is (80 GiB - 40 GiB or 80% - 50%) which is 30% or 40 GiB.

See these related values for more information Capacity Limit (%) and Adjusted Used Capacity (%).

This metric is not available for all storage systems, such as FlashSystem A9000, FlashSystem A9000R, and Dell EMC VMAX.

Used Capacity (GiB) and Used Capacity (%)

The amount of usable capacity that is taken up by the data in the pool after data reduction techniques have been applied.

Availability: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Recent Fill Rate (%)


The rate at which the capacity of the pool is being consumed over the last 30 days. Use this value to see how quickly your pools are filling up.

The Recent Fill Rate (%) tells you how quickly your pools are filling up

The Recent Fill Rate (%) of the pool is the difference between the fill rate % of the pool 30 days ago and today's fill rate % of the pool.

The fill rate % of the pool is calculated by dividing the capacity of the pool by its used capacity and multiplying it by 100.

Example



The fill rate % of the pool 30 days ago was $(40/100) \times 100$, which is 40%.

The fill rate % of the pool today is $(60/100) \times 100$, which is 60%.

So, in this example, the Recent Fill Rate (%) for the pool is $(60\% - 40\%)$, which is 20%.

If insufficient historical data is available to calculate the Recent Fill Rate (%), it is the difference between today's fill rate % and the oldest value for the fill rate % that can be calculated in the last 30 days.

Availability: This metric is not available for all storage systems, such as FlashSystem A9000, FlashSystem A9000R, and Dell EMC VMAX.

Recent Growth (GiB)

The growth in used capacity over the last 30 days. Use this value to identify the pools with the highest growth in storage consumption.

Recent growth is the difference between the current used capacity for the pool and the used capacity for the pool that was reported 30 days ago. If insufficient historical data is available, it is the difference between the current used capacity for the pool and the oldest value for the used capacity of the pool that was reported in the last 30 days.

Availability: This metric is not available for all storage systems, such as FlashSystem A9000, FlashSystem A9000R, and Dell EMC VMAX.

Tier Distribution (%)

For storage systems that support Easy Tier®, the distribution of the capacity of the volume extents across each drive class or tier.

Tip: To check the fill rate of all the pools in your storage systems, click [Storage > Block Storage Systems](#). To check the growth in used capacity, right-click any column heading on the Block Storage Systems page and click Recent Growth (GiB).

Monitor the total provisioned capacity of your volumes to identify the volumes that have sufficient capacity to meet your storage needs.

The following key capacity metrics provide information about storage usage and the availability of capacity for storing data:

Capacity (GiB)

The provisioned capacity of the volume.

Used Capacity (GiB) and Used Capacity (%)

The amount of provisioned capacity that is taken up by written data in a thin-provisioned volume after data reduction techniques are applied.

Unlike thin-provisioned volumes, which use capacity when it is needed, the capacity that is provisioned to standard-provisioned volumes is fully allocated and is no longer available to the pool. That is, the Used Capacity (%) of standard-provisioned volumes is reported as `Fully Allocated` and the Used Capacity (GiB) is the same value as the provisioned capacity.

Available for: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Available Capacity (GiB)

The difference between the provisioned capacity of the thin-provisioned volume and the used capacity of the thin-provisioned volume.

Because the pool's capacity is dedicated to the standard-provisioned volume, the Available Capacity (GiB) is reported as `0.00`.

Tier Distribution (%)

For storage systems that support Easy Tier®, the distribution of capacity for the volume extents across each drive class or tier.

Viewing information about enclosures

View information about enclosures. Enclosures are the metal structures in which various storage components are mounted, including control units, nodes, disks and drives, and power supplies.

Information about enclosures is available for storage systems, including systems that run IBM Spectrum Virtualize.

To view information about enclosures and their related resources, complete any of the following tasks:

Table 1. Viewing information about enclosures




Actions	Navigation
<p>View information about the enclosures that are associated with a storage system.</p> <p>Depending on the type of storage system, the following information about an enclosure is shown:</p> <ul style="list-style-type: none">• Status• Type of enclosure: Control, Expansion, or Storage• Machine Type model (MTM)• Physical location• I/O Group• The number of disks, drives, and slots that it contains• The number of nodes that it contains• The number of power supplies and canisters that it contains• The status of its power supplies (PSUs) and canisters	<ol style="list-style-type: none">1. In the menu bar, go to Storage >> Block Storage Systems.2. Right-click a storage system and select View Details.3. Under Internal Resources, click Enclosures.
<p>View the disks in an enclosure.</p>	<ol style="list-style-type: none">1. In the menu bar, go to Storage >> Block Storage Systems.2. Right-click a storage system and select View Details.3. Under Internal Resources, click Enclosures.4. In the Disks column, click the number to view more information about each of the disks.
<p>View the enclosure that a disk or drive is in.</p>	<ol style="list-style-type: none">1. In the menu bar, go to Storage >> Block Storage Systems.2. Right-click a storage system and select View Details.3. Under Internal Resources, click Disks or Drives.4. In the Enclosure column, view the name of the enclosure.5. (Optional) To view information about the enclosure, click its name.
<p>View the nodes in an enclosure.</p>	<ol style="list-style-type: none">1. In the menu bar, go to Storage >> Block Storage Systems.2. Right-click a storage system and select View Details.3. Under Internal Resources, click Enclosures.4. In the Nodes column, click the number to view more information about each of the nodes.
<p>View the enclosure that a node is in.</p>	<ol style="list-style-type: none">1. In the menu bar, go to Storage >> Block Storage Systems.2. Right-click a storage system and select View Details.3. Under Internal Resources, click Nodes.4. In the Enclosure column, view the name of the enclosure.5. (Optional) To view information about the enclosure, click its name.

Actions	Navigation
View the statuses of power supplies (PSUs) in an enclosure.	<ol style="list-style-type: none"> 1. In the menu bar, go to Storage > Block Storage Systems. 2. Right-click a storage system and select View Details. 3. Under Internal Resources, click Enclosures. 4. Right-click an enclosure and select View Properties. 5. View the Power Supply Status field.
View the statuses of canisters in an enclosure.	<ol style="list-style-type: none"> 1. In the menu bar, go to Storage > Block Storage Systems. 2. Right-click a storage system and select View Details. 3. Under Internal Resources, click Enclosures. 4. Right-click an enclosure and select View Properties. 5. View the Canister Status field.
<p>Add enclosures to general groups.</p> <p>Adding storage systems and their associated enclosures to general groups can be helpful when you want to receive alert notifications about changes to a group of logically related storage systems.</p>	<ol style="list-style-type: none"> 1. In the menu bar, go to Storage > Block Storage Systems. 2. Right-click a storage system and select View Details. 3. Under Internal Resources, click Enclosures. 4. Right-click an enclosure and select Add to General Group. 5. Add the enclosure to an existing group or create a new group and click Save.

Monitoring the status and condition of resources



Monitor the operational condition of storage systems, servers, hypervisors, fabrics, and switches and the status of their internal resources. Use this information to identify potential problem areas in a storage environment.








Table 1. Monitoring the status and condition of resources

	Explanation	Steps to view status or condition
Status	<p>The status of a resource that is reported by its own hardware to IBM Spectrum® Control. Statuses include Normal, Online, Offline, Degraded Ports, Operational, Error, Stopped, Starting, Completed, Unknown, and other statuses. Use the status to determine the condition, and if any actions must be taken to correct the problem. For example:</p> <ul style="list-style-type: none"> • If the cooling fans in a storage system are stopped and the internal temperature is too high, an error status  is reported by that storage system. • If a disk on a storage system is starting, a warning status  is reported for that disk by the storage system. <p>Tip: The status of internal resources is used to determine the operational condition of the associated top-level resources.</p>	<ol style="list-style-type: none"> 1. In the menu bar, go to the resource type that you want to view. For example, if you want to view the status of switches, go to Network > Switches. 2. Right-click a resource and select View Details. A status icon is shown next to the image of its related resource and its internal resources.
Condition	<p>The overall operational condition of a storage system, server, hypervisor, fabric, or switch. This condition represents the most critical status that was detected on the resource itself and on its internal resources.</p> <p>For example, if an error status was detected on a storage system pool, an error icon  is shown for the overall condition of the storage system. If no errors, warnings, or unreachable statuses were detected on a resource or on its internal resources, then a green symbol is shown for the condition of the storage system.</p> <p>Tip: Call Home events do not affect the condition of a top-level resource.</p>	<p>Dashboard view</p> <p>In the menu bar, go to Home > Dashboard. The icons that show overall condition are shown next to the image of a resource type.</p> <p>Resource list pages</p> <p>In the menu bar, go to the resource type that you want to view. For example, if you want to view the condition of switches, go to Network > Switches. The overall condition of a resource is displayed in the Condition column, and aggregated in the condition icons at the top left corner of its details page.</p>

IBM Spectrum Control IBM® Storage Insights provides a number of different icons to help you quickly determine the health of resources.

Table 2. Possible statuses and conditions of resources

Icon	Health	Explanation
	Error	A serious problem was detected on a resource or on its internal resources. Resolve these problems as soon as possible.
	Error - Acknowledged	<p>An Error status was detected and acknowledged. An Error - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.</p> <p>An acknowledged status is not used when determining the condition of related, higher-level resources. For example, if the status of a volume is Error, the condition of the associated storage system is also Error. If the Error status of the volume is acknowledged, its status is not used to determine the overall condition of the storage system.</p>

Icon	Health	Explanation
	Unreachable	A resource is not responding to requests from the IBM Spectrum Control server. This status might be caused by a problem in the network or by a Storage Resource agent that is no longer running and did not communicate that it was shutting down.
	Unreachable - Acknowledged	<p>An Unreachable status was detected and acknowledged. An Unreachable - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.</p> <p>An acknowledged status is not used when determining the condition of related, higher-level resources. For example, if the status of a controller is Unreachable, the condition of the associated server is also unreachable. If the Unreachable status of the controller is acknowledged, its status is not used to determine the overall condition of the server.</p>
	Warning	A Warning status represents potential problems on a resource or on its internal resources. This status is not critical.
	Warning - Acknowledged	<p>A Warning status was detected and acknowledged. A Warning - Acknowledged status indicates that a status was reviewed and is either resolved or can be ignored.</p> <p>A Warning - Acknowledged status is not used when determining the condition of related, higher-level resources. For example, if the status of a controller is Warning, the condition of the associated server is also Warning. If the Warning status of the controller is acknowledged, its status is not used to determine the overall condition of the server.</p>
	Normal	No warnings or errors were detected on a monitored resource.
	Not Monitored	For hosts, this status is displayed when IBM Spectrum Control monitors the storage system that the host is connected to, but the host itself was not added for monitoring. Unmonitored hosts are automatically created based on the host connections of monitored storage systems. Each host connection is represented as an unmonitored host.
	Unknown	A resource is known to IBM Spectrum Control but is not monitored. To change an Unknown status, run a probe to collect status information about the resource.

- [How the condition of a resource is determined](#)
The operational condition of a resource represents the most critical status that was detected on the resource itself and on its internal resources.
- [Viewing the overall condition of resources](#)
On the dashboard, you get an overview of the health of your storage resources such as fabrics, hypervisors, servers, storage systems, and switches and of storage consumers such as departments and applications.
- [Viewing the condition of specific types of resources](#)
Use resource list pages as a central location for viewing the condition of storage systems, servers, hypervisors, switches, and fabrics.
- [Viewing the status of resources](#)
Use detail pages to view the status of storage systems, servers, hypervisors, switches, and fabrics and the status of their internal resources. The status of a resource is reported to IBM Spectrum Control by its hardware.
- [Acknowledging the condition and status of resources](#)
Sometimes the status of resources might represent problems that commonly occur but can be ignored. In such cases, you can acknowledge those statuses so that they are not used when determining the overall condition of storage systems, servers, hypervisors, fabrics, and switches.
- [Monitoring vaults in IBM Cloud Object Storage](#)
Monitor access risk and storage risk for vaults in IBM Cloud Object Storage.


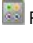


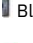
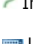




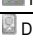





How the condition of a resource is determined






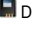
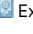

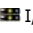
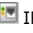












The operational condition of a resource represents the most critical status that was detected on the resource itself and on its internal resources.

To determine the condition of top-level resources (storage systems, servers, hypervisors, fabrics, and switches), IBM Spectrum® Control uses the status of those resources and the status of their internal resources.

The statuses of the following internal resources are used to calculate the overall condition of a top-level resource.

Table 1. Internal resources that are used to determine the condition of top-level resources

Top-level resource	Internal resources that are used to determine the condition of a top-level resource
 Fabric	 Ports  Switches
 Switch	 Blades  Inter-switch connections  Logical switches  Ports
	 Switches  Trunks  Blades  Ports
 Hypervisor	 Disks  Virtual machines  VMDKs

Top-level resource	Internal resources that are used to determine the condition of a top-level resource
 Server	 Controllers  Disks
 Block storage system	 Disks  Drives  External disks  FC ports  I/O groups  IP ports  Managed disks  Modules  Nodes  Pools  RAID arrays  Volumes
 File storage system	 Network shared disks  Nodes
 Object storage system	 Network shared disks  Nodes

The following statuses of internal resources are used to help calculate the condition of top-level resources:

- Normal
- Warning
- Error

Exceptions:






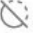





- Statuses that are acknowledged are not used to calculate the overall condition of a top-level resource.
- Events that are detected through Call Home do not affect the overall condition of a top-level resource. For example, a critical event is detected by Call Home, but the status of the storage system and its internal resources are normal based on metadata collection. In this case, the condition of the top-level resource remains as normal.

Internal resources for a top-level resource might have different statuses. IBM Spectrum Control[®] Storage Insights uses the most critical status of an internal resource to help determine the overall condition of a top-level resource. For example, in a storage system, a port might have an Error status, a pool might have a Warning status, and multiple controllers might have an Unknown status. In this case, if the storage system itself has a normal status, its overall condition is Error because it is the most critical status that was detected on internal resources.

The following table shows some of the possible combinations of statuses and the resulting, overall condition for a top-level resource.

Table 2. Propagation of the statuses for resources

Error 	Unreachable ¹ 	Warning 	Normal 	Unknown ² 	Resulting condition for a top-level resource
				X	 Unknown
			X		 Normal
			X	X	 Normal
		X			 Warning
		X		X	 Warning
		X	X	X	 Warning
	X				 Unreachable
	X			X	 Unreachable
	X		X	X	 Unreachable

Error 	Unreachable ¹ 	Warning 	Normal 	Unknown ² 	Resulting condition for a top-level resource
	X	X	X	X	 Unreachable
X					 Error
X				X	 Error
X			X	X	 Error
X		X	X	X	 Error
X	X	X	X	X	 Error

Note:

- The Unreachable status applies only to top-level resources.
- The Unknown status of an internal resource is not used to determine the condition of a top-level resource.

Viewing the overall condition of resources

On the dashboard, you get an overview of the health of your storage resources such as fabrics, hypervisors, servers, storage systems, and switches and of storage consumers such as departments and applications.

About this task

The condition of a resource represents the most critical status that was detected on that resource or on its internal resources.

Procedure

- From the Home menu, click Dashboard.
- Review the health of your storage environment.
The following conditions icons are shown for block storage systems:

-  Error
-  Unreachable
-  Warning
-  Normal

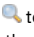
Viewing the condition of specific types of resources


Use resource list pages as a central location for viewing the condition of storage systems, servers, hypervisors, switches, and fabrics.

Procedure

- In the menu bar, go to the type of top-level resource that you want to view.
For example, if you want to view the condition of switches, go to Network > Switches.
- To view the total number of different conditions that were detected for the type of resource, check the condition icons on the resource list page.
The number next to the icon represents how many occurrences of each condition were detected.
For example, if 15 switches are monitored, but five switches have internal resources with a warning status, two switches have internal resources with an unreachable status, and three switches have internal resources with an error status, the following condition information is displayed:

-  5 Normal
-  5 Warning
-  2 Unreachable
-  3 Error

- To view the condition for specific resources, check the Condition column in the list of resources.
- Optional: Click the Search icon  to filter the list so that only resources with a specified condition are shown.
You can use filtering to identify the resources with the most serious conditions in your environment.
- Optional: Select Condition.
- Optional: In the Choose Value list, select the condition on which you want to filter.



The list of resources is automatically refreshed and only resources with the filtered status are shown. For example, select Warning  to view only the resources with a Warning condition.

Viewing the status of resources

Use detail pages to view the status of storage systems, servers, hypervisors, switches, and fabrics and the status of their internal resources. The status of a resource is reported to IBM Spectrum® Control by its hardware.

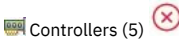
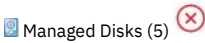
About this task


The status of a resource is different from the operational condition that is calculated by IBM Spectrum Control for top-level resources:

- *Condition* is determined by propagating the statuses of resources and is only shown for top-level resources.
- *Status* represents the status of a resource as reported by its hardware and is shown for top-level resources and internal resources. The following examples illustrate statuses that might be reported by a resource:
 - If the cooling fans in a storage system are stopped and the internal temperature is too high, an error status  is reported by that storage system.
 - If a disk on a storage system is starting, a warning status  is reported for that disk by the storage system.

Tip: View condition to identify which top-level resources might be encountering problems. View status to identify the specific resources that are causing the problems in a top-level resource.

Procedure

1. In the menu bar, go to the resource type that you want to view.
For example, if you want to view the status of block storage systems and their internal resources, go to Storage > Block Storage Systems.
2. Right-click a resource and select View Details.
3. To view the status of the top-level resource, view the status icon that is displayed next to its image on the page.
4. To view the aggregated status of an internal resource, view the status icon that is displayed next to its image in the Internal Resources section.
If a resource has a status other than normal, an icon is displayed for the most critical status.
For example, if a server has three controllers with an error status and two controllers with a warning status, the error status icon is shown for controllers in the Internal Resources section:

5. To view the status of a specific internal resource, click the name of that internal resource in the Internal Resources section.
A list of the monitored internal resources is displayed. For example, on the details page for a server, click Controllers (5) to view a list of the five controllers that are associated with the server.
6. Check the Status column to view the status for each internal resource.
Tip: When you view information about some internal resources, the Status column might show values that are more specific than Error, Unreachable, Warning, Unreachable, Unknown, and Normal. Use this additional status to determine the cause of the problem for an internal resource.
For example, in the Internal Resources section on the details page for a storage system, an error status might be shown for managed disks:


When you view the list of managed disks, the Status column might show an Offline status for a specific managed disk. When a managed disk is reported as offline, an error status icon  indicates that status.

Acknowledging the condition and status of resources

Sometimes the status of resources might represent problems that commonly occur but can be ignored. In such cases, you can acknowledge those statuses so that they are not used when determining the overall condition of storage systems, servers, hypervisors, fabrics, and switches.

- [Acknowledging the condition of top-level resources](#)
A summary of the conditions for storage systems, servers, hypervisors, fabrics, and switches is available in the Dashboard view and on resource list pages. You can acknowledge the condition of a top-level resource so that its condition is not shown as Error, Unreachable, or Warning on those pages.
- [Acknowledging the status of internal resources](#)
The statuses of internal resources are used to determine the condition of the associated storage systems, servers, hypervisors, fabrics, and switches. Sometimes these statuses might represent problems that commonly occur or are known, but can be ignored. In such cases, you can acknowledge the status of an internal resource so that it is not used to determine the condition of a top-level resource.

Acknowledging the condition of top-level resources

A summary of the conditions for storage systems, servers, hypervisors, fabrics, and switches is available in the Dashboard view and on resource list pages. You can acknowledge the condition of a top-level resource so that its condition is not shown as Error, Unreachable, or Warning on those pages.

Procedure

1. In the menu bar, go to the type of resource that has a condition you want to acknowledge.
For example, if you want to acknowledge the condition of a switch, go to Network > Switches.
2. In the list of resources, locate the resource with the Error, Unreachable, or Warning condition that you want to acknowledge.
3. Right-click the resource and select Mark Condition as Acknowledged.

A window shows that the condition for the specified resource was acknowledged.


- Optional: To acknowledge the condition of multiple resources at the same time, press Ctrl and click each resource. To select a series of resources, select the first resource, and then press Shift and click the last resource. When the resources are selected, right-click a row in the list and select Mark Condition as Acknowledged.
- Click Close.

Results

When you acknowledge the condition of a top-level resource, the following actions are taken:

Dashboard view

- The condition of the acknowledged resource is not used to determine the condition icon that is displayed for that resource type on the Dashboard view. For example, if you acknowledge the unreachable condition for a switch, the condition of that switch is not used to determine the number of unreachable conditions that is shown for switches on the Dashboard view.
- The acknowledged condition is shown as Acknowledged for a resource on the Dashboard view. For example, if 10 switches have Unreachable conditions and you acknowledge the condition for one of them, the following information is shown for switches:

 9 Unreachable (1 Acknowledged)

Resource details page

- The condition icons at the top of the page are updated. For example, if 10 switches have Unreachable conditions, and you acknowledge one of them, the total number Unreachable conditions decreases by one and the number of Unreachable - Acknowledged conditions increases by one. The following information is updated at the top of the Switches page:

 9 Unreachable  1 Unreachable - Acknowledged

- The value in the Condition column for the resource is updated to show an acknowledged condition. For example, if you acknowledge the Unreachable condition for a switch, the following value is shown in the Condition column:

 Unreachable

Related reference

- [Resources that you can monitor](#)
- [How the condition of a resource is determined](#)

Acknowledging the status of internal resources

The statuses of internal resources are used to determine the condition of the associated storage systems, servers, hypervisors, fabrics, and switches. Sometimes these statuses might represent problems that commonly occur or are known, but can be ignored. In such cases, you can acknowledge the status of an internal resource so that it is not used to determine the condition of a top-level resource.

About this task

For example, if the status of a volume is Error, the condition of the associated storage system is also Error. If the Error status of the volume is acknowledged, its status is not used to determine the overall condition of its storage system. In this case, if the other internal resources of the storage system are Normal, then the condition of the storage system is also Normal.

Procedure

- In the menu bar, go to the type of resource in which the internal resource is located.
For example, if you want to acknowledge the status of internal resources for a switch, go to Network > Switches.
- Right-click a resource and select View Details.
For example, if you want to acknowledge the status of a port on a switch, right-click that switch and select View Details.
- In the Internal Resources section of the resource details page, click the name of the internal resource.
For example, if you want to acknowledge the status of a port on a switch, click Ports.
- Right-click the resource and select Mark Condition as Acknowledged.
A window shows that the status for the specified resource was acknowledged. For example, if you want to acknowledge the status of a port, right-click the port and select Mark Condition as Acknowledged.
- Optional: To acknowledge multiple resources at the same time, press Ctrl and click each resource. To select a series of resources, select the first resource, and then press Shift and click the last resource. When the resources are selected, right-click a row in the list and select Mark Condition as Acknowledged.
- Click Close.
The status of the related, top-level resource might take several minutes to update.

Results

When you acknowledge the status of an internal resource, the following actions occur:

- The status of the internal resource is no longer used to determine the condition of the associated top-level resource.
- The status icons at the top of the page for the internal resource are updated.
For example, if you acknowledge the Error status for a controller on a server, the total number of Error statuses decreases by one and the number of Error - Acknowledged statuses increases by one on the Controllers page. If originally there were five controllers with Error statuses, the following updated status information is shown:

Related reference

- [Resources that you can monitor](#)
- [How the condition of a resource is determined](#)

Monitoring vaults in IBM Cloud Object Storage

Monitor access risk and storage risk for vaults in IBM® Cloud Object Storage.

You can use this information to complete the following tasks:

- Identify the vaults that cannot be accessed and the vaults that are at risk of access failure.
- Identify the vaults that are not available for write operations and the vaults that are at risk of storage failure.
- [Monitoring access to vaults in IBM Cloud Object Storage](#)
In these scenarios, you use IBM Spectrum® Control to identify the vaults that cannot be accessed and the vaults that are at risk of access failure. Then, you identify the access pools and the COS Accesser® nodes that the vaults depend on.
- [Calculating the failure tolerance for vaults](#)
The failure tolerance value for vaults represents the tolerance of a vault to drive failures across all of the COS Slicestor® nodes in its storage pool. The lower the value for drive failure tolerance, the greater the risk that write operations can no longer be completed for the vault.

Monitoring access to vaults in IBM Cloud Object Storage

In these scenarios, you use IBM Spectrum® Control to identify the vaults that cannot be accessed and the vaults that are at risk of access failure. Then, you identify the access pools and the COS Accesser® nodes that the vaults depend on.

How the accessibility value for a vault is determined

The accessibility value that is calculated for each vault comprises the number of COS Accesser nodes that are available to access the vault. For example, if five COS Accesser nodes are configured for a vault and four of the nodes are in a failed state, then the accessibility value is one.

On the details page for IBM® Cloud Object Storage, the Vaults by Accessers chart shows the vaults that are inaccessible and the accessibility values for the remaining vaults in the storage system.

- [Investigating vaults that cannot be accessed](#)
In this scenario, you use IBM Spectrum Control to identify the vaults that cannot be accessed because of COS Accesser node failures.
- [Monitoring vaults that are at risk of access failure](#)
In this scenario, you use IBM Spectrum Control to check which vaults are at risk because of COS Accesser node failures.

Investigating vaults that cannot be accessed


In this scenario, you use IBM Spectrum® Control to identify the vaults that cannot be accessed because of COS Accesser® node failures.

About this task

You have a backup application that uses IBM® Cloud Object Storage vaults. Because the application is critical to your business, you want to quickly identify vaults that cannot be accessed because of COS Accesser node failures. You also want to know the access pools that the vaults depend on and the COS Accesser nodes that failed in those pools.

Procedure

1. To determine whether any of your vaults cannot be accessed, click Storage > Object Storage Systems.
2. Right-click the storage system that is used by your backup application and click View Details.
3. In the navigation pane, click Overview.
The first line of the Vaults by Accessers chart shows the number of vaults that cannot be accessed, for example, Inaccessible 4.
4. To identify the vaults that are inaccessible, click Vaults in the navigation pane. To sort the data, click the Accessibility column.

A value of 0 in the Accessibility column indicates that the vault is inaccessible. The error status  indicates that the vault is inaccessible because none of the COS Accesser nodes that are configured for the vault are available. In this scenario, all the inaccessible vaults are configured to use just one access pool. The name of that access pool is shown in Access Pools.



Vaults

✓ 104 Normal
⚠ 2 Warning
✗ 5 Error

Actions		
Name	Access Pools	Accessibility
accessertest1	AccessPool2	✗ 0
demovault7	AccessPool2	✗ 0
myvault	AccessPool2	✗ 0
vault2	AccessPool2	✗ 0

5. To identify the COS Accesser nodes that failed in the access pool, click the access pool name. Then, click the Accesser Nodes tab.

What to do next

To restore access to the vaults, fix the COS Accesser nodes that failed or add COS Accesser nodes to the access pool.

Monitoring vaults that are at risk of access failure

In this scenario, you use IBM Spectrum® Control to check which vaults are at risk because of COS Accesser® node failures.

About this task

You have a backup application that uses IBM® Cloud Object Storage vaults. Because the application is critical to your business, you want to monitor the tolerance of the vaults to COS Accesser node failures and identify the vaults that are at risk of access failure.

For the vaults with a high risk of failure, you want to identify the access pools that the vaults depend on and the status of the COS Accesser nodes in the access pools.

Procedure

1. Click **Storage > Object Storage Systems**.
2. Right-click the storage system that is used by your backup application and click **View Details**.
3. In the navigation pane, click **Overview**.

The **Vaults by Accessers** chart shows the tolerance of the vaults to COS Accesser node failures. For example, a line on the chart that shows a value of 1 **Node Failure 57 Vaults at Risk** indicates that 57 vaults rely on a single COS Accesser node for access. If the node fails, the vault cannot be accessed.

4. To identify the vaults with an accessibility value of 1, click **Vaults** in the navigation pane. To sort the data, click the **Accessibility** column.



Vaults

✓ 104 Normal
⚠ 2 Warning
✗ 5 Error

Actions		
Name	Access Pools	Accessibility
failovertest	2	⚠ 1
test2	2	⚠ 1
DemoVault	AccessPool1	1
arun878172498111...	AccessPool1	1
arunagiri123	AccessPool1	1
arunagiri123.meta	AccessPool1	1
containermodetest	AccessPool1	1


In this scenario, there are some vaults that depend on only one access pool. The name of that access pool is shown in Access Pools. For other vaults that depend on multiple pools, a number is shown in Access Pools instead.

Tip: To reduce the risk of access failure for the vaults that depend on only one access pool, add COS Accesser nodes to the access pool or deploy the vaults to other access pools.

For the vaults that depend on multiple access pools, the warning status  indicates that some of the COS Accesser nodes in the pools are not available.

5. To investigate the cause of the access warning, click the number in the Access Pools column.

You see that one of the access pools is shown with an error status , which indicates that the COS Accesser nodes in the pool are not available.



6. To identify the COS Accesser nodes that failed in the access pool, click Accesser Nodes in the navigation pane. To sort the data, click the Access Pool column.

What to do next

Fix the COS Accesser nodes that failed or add COS Accesser nodes to the access pool.

Calculating the failure tolerance for vaults

The failure tolerance value for vaults represents the tolerance of a vault to drive failures across all of the COS Slicestor® nodes in its storage pool. The lower the value for drive failure tolerance, the greater the risk that write operations can no longer be completed for the vault.

The calculation of drive failure tolerance is based on the following values:

- The write threshold value in the IDA that is configured for the vault. The write threshold is the minimum number of COS Slicestor nodes that must be available to complete write operations for the vault.
- The number of available COS Slicestor nodes in the storage pool that the vault belongs to.
- The value for drive failure tolerance for each of the available COS Slicestor nodes in the storage pool. This value represents the tolerance of the node to drive failures and is based on the number of failed drives in the node and the drive error threshold. A value of 2 means that the node will be unavailable to store vault data if 2 drives fail in the node.

For example, a vault in your IBM® Cloud Object Storage environment is configured with an IDA of 8-4-6. The IDA determines that the vault data is stored across 8 COS Slicestor nodes and a minimum of 6 nodes must be available to complete write operations for the vault.

The nodes have the following values for drive failure tolerance:

Node	Drive Failure Tolerance
node_1	3
node_2	4
node_3	0*
node_4	2
node_5	4
node_6	1
node_7	3
node_8	3

*A failure tolerance value of 0 means that node_3 is not available to store vault data.

7 COS Slicestor nodes are available in the storage pool and a minimum of 6 nodes must be available to complete write operations. Therefore, if 2 more nodes fail, the write threshold value is exceeded and write operations can no longer be completed for the vault.

Use the following steps to calculate the value for drive failure tolerance for a vault:

1. Use the IDA column on the Vaults page to identify the write threshold value for the vault. In the example, the write threshold value is 6.
2. Use the Slicestor Nodes page to determine the number of available COS Slicestor nodes in the storage pool that the vault belongs to. A value > 0 in the Drive Failure Tolerance column means that the node is available. In the example, 7 nodes are available.
3. Calculate how many nodes must fail for the number of available nodes to fall below the write threshold value. In the example, $(7 - 6) + 1 = 2$. When 2 nodes fail, the vault is not available for write operations.
4. Use the Drive Failure Tolerance column on the Slicestor Nodes page to identify the 2 available nodes in the storage pool that have the lowest tolerance to drive failures.

In the example, the 2 nodes are as follows:

Node	Drive Failure Tolerance
node_4	2
node_6	1

5. To calculate the overall failure tolerance for the vault, sum the failure tolerance values for the 2 nodes. $2 + 1 = 3$.
The drive failure tolerance for the vault is 3. If 3 drives fail, write operations can no longer be completed for the vault.

Monitoring the performance of resources

IBM Spectrum® Control can collect information about the performance of storage systems and switches. This information includes key performance metrics that can help you measure, identify, and troubleshoot performance issues and bottlenecks in your storage.


You can use the performance monitoring to complete the following tasks:

- Measure, compare, and troubleshoot the performance of switches, storage systems, and their internal resources.
- Review the alerts that were triggered when the performance of a resource fell outside of a specific range.
- Customize views of performance so that you can analyze specific resources and metrics during time ranges that you specify.

- View performance information in a chart or table format to help you quickly identify where and when performance issues are occurring. The chart is a visual representation of how the performance of resources trend over time.
- Drill down into resources to view detailed information about the performance of internal and related resources. For example, if a SAN Volume Controller is shown in the chart, you can quickly view and compare the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server, IBM Spectrum Control correlates the server with the known host connections on monitored storage systems. If matches are found, you can view the amount of storage that is assigned to the server, and trace that storage back to the storage system. You can then view details about the internal resources of the related storage system, including performance information.
- Export performance information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.

Before you begin

Before you can view performance information for resources, you must complete the following tasks:

Task	 Learn more
Add storage systems and switches for monitoring by IBM Spectrum Control.	For information about how to add these resources, see the following topics: <ul style="list-style-type: none"> • Adding storage systems • Adding fabrics and switches
Ensure that performance data is being collected for a resource. Typically, data collection is scheduled when resources are added for monitoring.	For information about verifying data collection and scheduling performance monitors, see the following topics: <ul style="list-style-type: none"> • Verifying that a performance monitor is running for a resource • Creating performance monitors in IBM Spectrum Control
Define performance alerts to be notified if the performance of a resource falls outside a specified range and might represent a potential problem. When you define a performance alert for a resource, select a specific metric that you want to measure. For example, you can define a threshold that notifies you when the total I/O port rate for a storage system falls outside a specified range.	For information about how to define performance alerts, see Defining alert definitions for performance changes .

- [Viewing performance information](#)
View the performance of storage systems, switches, and their internal resources. You can also view the performance of resources that violated a performance threshold and generated an alert.
- [IBM Spectrum Virtualize guideline values for key performance indicators](#)
To improve the performance and resiliency of your storage environment, compare the guideline values for key performance indicators with the values reported for your storage systems and devices.
- [Viewing performance alerts](#)
View the alerts that were generated when the measured value of a performance metric falls outside of the specified threshold for a resource.
- [Exporting performance data for storage systems and fabrics](#)
To help resolve performance issues with storage systems and fabrics, you can export performance data for the resources to a compressed file. If you contact IBM® Support to help you analyze the problem, you might be asked to send this file.
- [Performance metrics](#)
IBM Spectrum Control can collect information about the performance of storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks. You can view performance information or specify alerts for when the performance of a resource falls outside a specified range.

Related tasks

- [Tutorial: Identifying the source of slow drain problems caused by depletion of buffer credits](#)

Related reference

- [Performance information is not displayed for a resource](#)

Viewing performance information



View the performance of storage systems, switches, and their internal resources. You can also view the performance of resources that violated a performance threshold and generated an alert.

Before you begin

Before you view the performance of resources, ensure that performance data was collected for those resources during the time ranges that you want to analyze. IBM Spectrum® Control uses performance monitors to collect metrics for measuring the performance of storage systems and switches. For information about how to verify that performance monitors were run, see [Verifying that a performance monitor is running for a resource](#).

About this task

You can access the performance view for a resource from different locations in IBM Spectrum Control:

Task	Accessing the performance view in the GUI
View the performance of specific storage systems, switches, and their internal resources.	<ol style="list-style-type: none"> 1. In the menu bar, go to the type of resource that you want to view. For storage systems, depending on the type of storage system that you want to view, go to Storage > Block Storage Systems, Storage > File Storage Systems, or Storage > Object Storage Systems. For switches, go to Network > Switches. 2. Right-click a resource and select View Performance. By default, the following information is displayed in a chart for each type of resource: <ul style="list-style-type: none"> • For storage systems, the top 5 storage systems with the highest total I/O rate are shown. • For switches, the top 5 switches with the highest total port data rate are shown. 3. To view the performance of an internal resource, right-click the top-level resource in the chart legend and select the option for the internal resource that you want to view. For example, for a SAN Volume Controller, right-click its name in the chart legend and select MDisk Performance to view the performance of its managed disks. For managed disks, the top 5 contributors to the storage system workload are shown on the chart. The default metrics are total back-end I/O rate and total back-end response time. <p>Tips:</p> <ul style="list-style-type: none"> • To add or change the metrics that are displayed for a resource, click the select metrics icon  and select the metrics to display. • You can also view the performance of internal resources on the details page for the related, top-level resource. For more information, see Table 1. • To view the performance of multiple resources at the same time, press Shift or Ctrl and click the resources that you want to view. Press Shift and click to select consecutive rows in the list of resources; press Ctrl and click to select non-consecutive rows. Then, right-click any of the selected resources and select View Performance. Each resource is represented by a separate line in the performance chart and separate rows in the table view.
View the overall performance of resources of the same type.	<ol style="list-style-type: none"> 1. In the menu bar, go to the type of resource that you want to view. For storage systems, depending on the type of storage system that you want to view, go to Storage > Block Storage Systems, Storage > File Storage Systems, or Storage > Object Storage Systems. For switches, go to Network > Switches. 2. Click the Performance tab. By default, the following information is displayed in a chart: <ul style="list-style-type: none"> • For storage systems, the top 5 storage systems with the highest total I/O rate are shown. • For switches, the top 5 switches with the highest total port data rate are shown. <p>Tip: To add or change the metrics that are displayed for a resource, click the select metrics icon  and select the metrics to display.</p>
View the alerts that were generated by threshold violations and the performance of the resources where the violation occurred.	For information about how to view performance alerts, see Viewing performance alerts .
View the performance of a resource that is included in a specific performance monitor.	<ol style="list-style-type: none"> 1. In the menu bar in the web-based GUI, go to Home > Performance Monitors. 2. In the Name column, locate the name of the resource that you want to view. 3. Right-click the performance monitor and select View Performance.
View predefined reports and create custom reports that show the performance of resources. In the optional Cognos® Analytics reporting tool, you can view reports about storage systems, components of storage systems, switches, and switch ports.	<ol style="list-style-type: none"> 1. Go to the URL for your Cognos Analytics server. The URL is similar to this URL: http://myhostname:9300/bi 2. Click Team Content in the Welcome portal. 3. To view a predefined performance report, complete the following steps: <ol style="list-style-type: none"> a. Click IBM Spectrum Control Predefined Reports. b. Navigate to the performance report that you want to view and click the name of the report. 4. To create custom performance reports, complete the following steps: <ol style="list-style-type: none"> a. Click IBM Spectrum Control Packages. b. Right-click Performance, then click Create report. c. Click a template for the report.

You can also view the performance of internal resources on the details pages for the related, top-level resources. For example, you can view the performance of the controllers for a storage system on the details page for that storage system. To access the performance view for an internal resource on a details page, select the type of internal resource and click the Performance tab.

The following information summarizes the internal resources and locations where you can access a performance view.

Table 1. Internal resources and locations where you access the performance view

Internal resource	Locations in the GUI where you can access the performance view
Pools	<ul style="list-style-type: none"> • Storage > Pools • Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Pools <p>Available only for pools that are associated with the following storage systems and virtualizers: DS8000®, SAN Volume Controller, the IBM® Storwize® family, IBM Spectrum Accelerate, and XIV® systems.</p>
Volumes	<ul style="list-style-type: none"> • Storage > Volumes • Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Volumes
Controllers	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Controllers
Disks	<p>Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Disks</p> <p>Available only for disks that are associated with storage virtualizers such as SAN Volume Controller and the IBM Storwize family.</p>

Internal resource	Locations in the GUI where you can access the performance view
Host Connections	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Host Connections Available only for host connections that are associated with the following storage systems and virtualizers: DS8000, SAN Volume Controller, the IBM Storwize family, and XIV systems.
I/O Group	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click I/O Groups
Modules	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Modules
Nodes	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Nodes
Managed Disks	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Managed Disks Available only for nodes that are associated with storage virtualizers such as SAN Volume Controller, the IBM Storwize family, Hitachi VSP, and NetApp storage systems running ONTAP 9.
Ports on storage systems	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click Ports
Ports on switches	Network > Switches, and click Ports
RAID Arrays	Click Storage, and then click Block Storage Systems, File Storage Systems, or Object Storage Systems, depending on the storage system that you require. Right-click the storage system, click View Details, and click RAID Arrays Available only for RAID arrays that are associated with DS8000 storage systems.

- [How performance information is displayed](#)
The performance view is displayed when you view the performance of resources, which includes resources in alerts and performance monitors. Information in the performance view is organized into two main sections: a chart or table and a legend.
- [Saving URLs for performance views](#)
You can create a bookmark or a favorite for a performance view.
- [Identifying performance issues for IBM Spectrum Virtualize storage systems](#)
You can identify key metrics that are outside of a standard range for resources that run IBM Spectrum Virtualize by using the key performance indicators charts.

Related information

- [Performance statistics \(for SAN Volume Controller systems\)](#)
- [Performance statistics \(for Storwize V7000 Unified systems\)](#)

How performance information is displayed

The performance view is displayed when you view the performance of resources, which includes resources in alerts and performance monitors. Information in the performance view is organized into two main sections: a chart or table and a legend.

Performance chart

The top section of the performance view shows information about the selected resources. You can view this information in the following formats:

Chart

The chart shows a visual representation of how the performance of a resource trends over time. Each line on the chart represents a metric and a resource. For example, if you select two metrics and three resources, six lines are shown on the chart. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

Table

The table shows performance and asset information that is formatted into rows and columns. Each row represents a resource and a time stamp from the chart; each column represents a metric from the chart. For example, if two metrics and three resources are displayed on a chart, and each of the lines on the chart have 10 data points, 30 rows and 2 columns are shown in the table.
To view other metrics and asset information for a resource, right-click anywhere in the header row of the table and select extra columns. The type of resource determines the metrics and information that is available.

Performance chart legend

The bottom section of the performance view shows more information about the resources in the chart. The information is formatted into rows and columns. Each row represents a resource that was selected for the view. Each column provides asset and performance information about a resource.

Each metric on the chart is represented by two columns in the legend. The first column includes the weighted average of the metric for each resource in the selected time range. The second column represents the peak data point of the metric for each resource in the selected time range. For example, if a metric in the chart is Total I/O Rate, the following columns are displayed in the legend: Avg Total I/O Rate, Worst Total I/O Rate. The worst value can be the minimum or the maximum number, depending on the metric. For example, the worst value for an I/O rate metric is a maximum number, while the worst value for the Cache Hit Percentage metric is a minimum number.

You can also view information about the resources that are related to the resources in the legend. For example, if a SAN Volume Controller is listed in the legend, right-click it to view the performance of its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage. For more information about the resources that you can view in the chart legend, see [Resources in the chart legend](#).

Tip:

- If you hide a resource in the performance chart, the row for that resource remains visible in the chart legend.
- In the chart legend, information about a resource represents the state of that resource during the selected time range. Specifically, this information does not represent the current state of a resource, but instead shows the final state of the resource during the selected time range.

For example, if the current date is January 1, and you view a volume with the time range set to December 1 to December 7, the legend shows the capacity of that volume on December 7. If the capacity of the volume was changed between December 7 and January 1, this historical capacity is different from the current capacity of the volume on January 1.

- **Controls**
Each performance view includes controls for customizing how information is displayed. The type of view and the resources that you are viewing determine which controls are available.
- **Resources in the chart legend**
In the performance view, performance metrics and related information for resources are shown in a chart and in the chart legend. For resources in the chart legend, you can open separate performance views for their internal or related resources. The separate performance views use the same time range as the performance view in the main window of the GUI.

Controls

Each performance view includes controls for customizing how information is displayed. The type of view and the resources that you are viewing determine which controls are available.



View chart

View performance information in a chart format. The chart shows a visual representation of how resource performance trends over time. Each line on the chart represents a metric and a resource. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

Hover the mouse pointer over points on a line to view a snapshot of performance information at a specific time.



View table

View performance information in a table format. Each row represents a resource and a time stamp. Each column represents a metric. You can view other metrics and information for a resource by right-clicking anywhere in the header row for the table and selecting more columns. The type of resource determines the metrics and information that is available.

For information how to filter, sort, and customize the columns in a table, see [Customizing lists](#).



Metrics

Add metrics to the performance chart. On the Select Chart Metrics dialog, metrics are organized into Volume, Disk, Pool, Port, and Node categories. The metrics and categories that are available depend on the type of resource that is being shown in the chart. The number next to the name of the category represents the number of metrics that are currently selected from that category.

You can select multiple metrics at the same time and from different categories, but you cannot include more than two unit types in the same view. For example, if you select metrics that use % and ops/s as units of measurement, you cannot select more metrics that use different units of measurement such as KiB/op or MiB/s.

For a list of metrics that you can view for resources, see [Performance metrics](#).



Hide and show resources

When the chart includes multiple resources, you can click the icon next to a resource to show only the line for that resource. Each icon is shown in a different color to match the color of line for the resource.

You can also show and hide resources in the chart by selecting resources in the chart legend. To select multiple resources at the same time, press Shift or Ctrl and click those resources. Press Shift and click to select consecutive rows in the legend; press Ctrl and click to select non-consecutive rows.

You can show up to 10 resources in a chart at the same time.



Specify granularity

Determine the granularity of the data that is shown in a performance view. Granularity determines the points that are shown on the y-axis of a chart, the points on the lines in a chart, and the rows in a table. When you first view performance information, the default time range is the last 12 hours and the granularity matches the frequency of data collection by the performance monitor (sample).

To change the granularity of chart or table, select one of the following options from the granularity menu:



(daily)

Set the granularity to show 1-day increments for performance information that is shown on the chart and table.

This granularity is only available if the time range of the chart is more than 1 day.



(hourly)

Set the granularity to show one-hour increments for performance information that is shown on the chart and table.

This granularity is only available if the time range of the chart is more than 1 hour.



(Sample)

Set the granularity to match the frequency of data collection by a performance monitor. For example, if a performance monitor is scheduled to collect sample data every 15 minutes, each point on the chart is shown in 15-minute increments.

This granularity is only available if the time range of the chart is less than 2 weeks.

If the performance monitor collects sample data at 1-minute intervals, data is displayed as 1-minute intervals in the chart only when data is available for all of the selected time range. If data is not available at 1-minute intervals for any portion of the selected time range, the data is displayed in 5-minute intervals. For example, the performance monitor for a resource collects sample data every 5 minutes. Six hours ago, the performance monitor was changed to start collecting data every minute. You want to view sample data for the resource for the last 12 hours. However, the performance monitor did not collect sample data every minute for 6 of the last 12 hours. Therefore, the data is displayed in 5-minute intervals rather than 1-minute intervals.

Restriction: Each granularity has a maximum time range for displaying data. This time range is dynamically calculated based on the resources and metrics that are shown on the chart, and the data points per line. When the maximum range is exceeded, the granularity icon is disabled and a message indicates that there are too many lines on the chart. To avoid this problem, reduce the time range or the number of resources or metrics that are selected.

Specify a time range

The time range of the performance information is shown below the chart. You can change this range to display information for different times when data was collected. When you first access the performance view, the default time range is the last 12 hours.

Ensure that performance data was collected during the time range that you select. If data was not collected during the time range, the chart and table are blank. If data collection was interrupted during the time range, the chart and table show gaps for the time increments when data was not collected.

For example, if you select a time range for the last 7 days, but data was not collected on days 4 and 5, the lines in the chart do not show data for days 4 and 5.



Hide and show chart controls

The performance view lists the names of resources that are being shown in the chart and the controls for customizing that chart. Click the left arrow icon to hide the controls and display the chart in the entire window. Click the right arrow icon again to show the controls.



Export information about the chart to a file

Export information on a performance view to a CSV file.



Open the performance view in a separate web browser window

Open a duplicate of the current performance view in a separate web browser window. You can change the information that is displayed in this separate window while retaining the original performance view for comparison.



Synchronize the time range across all the open performance views

Synchronize the time range across all the performance views that are displayed in separate browser windows. Use this action when you change the time range in a performance view and want to apply the same time range to the other performance views.

Restriction: This action affects only the performance views that are displayed in separate browser windows. The time range of the performance view that is shown as part of the main window for the GUI is not affected.

For example, you can view the storage system volumes that are assigned to a server. You can then open separate browser windows for the host connections, pools, and managed disks that are related to one of the volumes. If you change the time range in the performance view for host connections, click the synchronize button to apply the same time range to the views for pools and managed disks. The performance view of the original volume is not changed.

Actions for managing the resources in the chart legend

The chart legend in bottom section of the view shows more information about the selected resources. This information is organized into rows and columns, where each row represents a resource.

When you select one or more resources in the legend, the following actions are available in the Actions menu:

View Properties

View key details about a resource, including asset, status, configuration, capacity, and performance information.

View *resource* Performance

View the performance of resources that are internal or related to a resource in the chart legend. For example, when you view the performance of a SAN Volume Controller, you can right-click it and view the performance of its internal resources, such as disks, volumes, ports, managed disks, and back-end storage.

Information about an internal or related resource is shown in a separate web browser window. This window uses the same time range as the performance view in the main window of the GUI.

Restriction: If performance data was not collected for an internal or related resource, that resource is not shown in the view. For example, if you right-click a storage system and select Volume performance, only the volumes for which performance data was collected are shown.

For more information about the resources that you can view in the chart legend, see [Resources in the chart legend](#).

Key Metrics View

View a set of predefined metrics for resources in the chart legend.

Custom View

Customize the set of metrics that are shown for resources in the chart legend.

Resources in the chart legend

In the performance view, performance metrics and related information for resources are shown in a chart and in the chart legend. For resources in the chart legend, you can open separate performance views for their internal or related resources. The separate performance views use the same time range as the performance view in the main window of the GUI.

For example, when you view the performance of a SAN Volume Controller, you can right-click it in the chart legend to open a performance view for its internal and related resources, such as disks, volumes, ports, managed disks, and back-end storage.

The type of resource determines the internal and related resources that you can view. The following tables show the internal and related resources that are available when you right-click a resource in the chart legend.

Table 1. Resources that you can view for DS8000 storage systems in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Volume, RAID array, pool, controller, host connection, port	Storage virtualizer, switch
Volume	Storage system, RAID array, pool, controller, host connection	Managed disk
RAID array	Storage system, volume, pool, controller, host connection	None
Pool	Storage system, volume, RAID array, controller, host connection	None
Controller	Storage system, volume, RAID array, pool, host connection	None
Host connection	Storage system, volume, RAID array, pool, controller	None
Storage system port	Storage system	Switch port

Table 2. Resources that you can view for SAN Volume Controller, Storwize V3500, Storwize V3700, Storwize V7000, and Storwize V7000 Unified storage systems in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Volume, pool, node, I/O group, disk, managed disk, host connection, port	Back-end storage system, storage virtualizer, switch
Volume	Storage system, pool, node, I/O group, disk, managed disk, host connection	None
Pool	Storage system, volume, node, I/O group, disk, managed disk, host connection	None
Node	Storage system, volume, pool, I/O group, disk, managed disk, host connection, port	None

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
I/O group	Storage system, volume, pool, node, disk, managed disk, host connection, port	None
Disk	Storage system, volume, pool, I/O group, node, managed disk, host connection	None
Managed disk	Storage system, volume, pool, node, I/O group, disk, host connection	Back-end volume
Host connection	Storage system, volume, pool, node, I/O group, disk, host connection	None
Storage system port	Storage system, I/O group, node	Switch port

Table 3. Resources that you can view for IBM Spectrum Scale in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	File system, node	None
File system	Storage system, node	None
Node	Storage system, file system	None

Table 4. Resources that you can view for IBM Spectrum Accelerate in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Module, pool, volume, host connection, port	Storage virtualizer, switch
Module	Storage system, pool, volume, host connection, port	None
Pool	Storage system, module, volume, host connection	None
Volume	Storage system, module, pool, host connection	Managed Disk

Table 5. Resources that you can view for XIV systems in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Module, pool, volume, host connection, port	Storage virtualizer, switch
Module	Storage system, pool, volume, host connection, port	None
Pool	Storage system, module, volume, host connection	None
Volume	Storage system, module, pool, host connection	Managed Disk
Host connection	Storage system, module, pool, volume	None
Storage system port	Storage system, module	Switch port

Table 6. Resources that you can view for IBM Spectrum Accelerate in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Module, pool, volume, host connection, port	Storage virtualizer, switch
Module	Storage system, pool, volume, host connection, port	None
Pool	Storage system, module, volume, host connection	None
Volume	Storage system, module, pool, host connection	Managed Disk

Table 7. Resources that you can view for SMI-S BSP systems in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Storage system	Volume, port	Storage virtualizer, switch
Volume	Storage system	Managed disk
Storage system port	Storage system	Switch port

Table 8. Resources that you can view for switches in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Switch	Port	Storage system, switch
Port	Switch	Storage system port, switch port

Table 9. Resources that you can view for servers in the chart legend

Resource in the chart legend	Right-click the resource in the chart legend to view the performance of the following resources	Right-click the resource in the chart legend to view the performance of the following related resources
Server	None	Storage system, volume, pool, controller, module, node, I/O group, disk, managed disk, host connection, switch, storage system port, switch port

Saving URLs for performance views

You can create a bookmark or a favorite for a performance view.

About this task

The link in the bookmark or favorite stores details of the performance view, such as the resources, metrics, interval, and other details. The link also stores the time period that is specified in the time selector at the top of the chart.

You can also copy the URL from the address bar if you want to share the URL with colleagues.

Procedure

Do one of the following steps:

- In Firefox, right-click in the performance view and click Bookmark This Page.
- In Internet Explorer, right-click in the performance view and click Add to favorites.

Identifying performance issues for IBM Spectrum Virtualize storage systems

You can identify key metrics that are outside of a standard range for resources that run IBM Spectrum Virtualize by using the key performance indicators charts.

Before you begin

Before you view the performance of your resources that run IBM Spectrum Virtualize, ensure that performance data was collected for those systems during the time ranges that you want to analyze. IBM Spectrum® Control uses performance monitors to collect data so that you can measure the performance of storage systems and switches. To verify that performance monitors were run, follow the instructions in [Verifying that a performance monitor is running for a resource](#).

About this task

In this documentation, IBM Spectrum Virtualize is used to refer collectively to IBM® SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, and IBM Storwize® storage systems, and to IBM FlashSystem® devices that run IBM Spectrum Virtualize.

Restriction: Key performance indicators are provided only for IBM FlashSystem 9100 from the IBM FlashSystem family. Key performance indicators are not provided for file storage in Storwize V7000 Unified storage systems.

Procedure

1. To select a storage system, choose one of the following options from the menu bar:

Resource	Menu bar
Block storage systems	Storage > Block Storage Systems
File storage systems	Storage > File Storage Systems

2. From the list of storage systems that are displayed, right-click the resource that you want to view and click View Details.
For example, select a SAN Volume Controller or a Storwize storage system on the Block Storage Systems page, or select a Storwize V7000 Unified storage system on the File Storage Systems page.
3. In the General section of the resource details page, click Performance.
The key performance indicators charts are displayed with data that was collected from monitored resources over the last 24 hours. The performance of each monitored resource is charted against the best practice guidelines for IBM Spectrum Virtualize storage systems.
4. Optional: If you want to view the performance charts for a period other than the default 24-hour period, change the date for all of the charts. Click Last 24 hours and use the calendar to specify a date.
5. Optional: Drill down to view detailed information about a specific I/O group, port, or node. Click a resource from the chart's list.
For example, select the name of an I/O group in the Read Response Time by I/O Group chart to view the average number of milliseconds taken by the I/O group to complete a read operation.

By default, metrics are shown on the charts for the following time periods:

- The last 24 hours
- The same 24-hour period for yesterday
- The same 24-hour period for the day seven days before today

If you selected a date from the calendar, the metrics are shown on the charts for the following time periods:

- The selected date
- The day before the selected date
- The day seven days before the selected date

IBM Spectrum Virtualize guideline values for key performance indicators

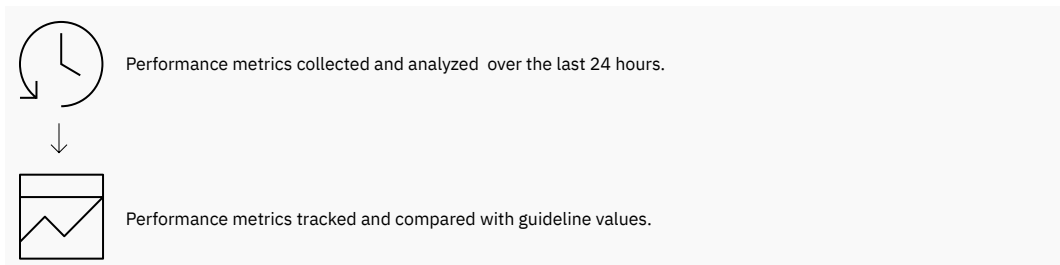
To improve the performance and resiliency of your storage environment, compare the guideline values for key performance indicators with the values reported for your storage systems and devices.

Guideline values for performance were established by monitoring, measuring, analyzing, and stress testing the performance of IBM Spectrum Virtualize storage systems.

The name IBM Spectrum Virtualize is used to refer to the following types of storage systems:

- IBM® SAN Volume Controller
- IBM Spectrum Virtualize for Public Cloud
- IBM Spectrum Virtualize as Software Only
- IBM Storwize® block storage systems
- IBM FlashSystem® devices that run IBM Spectrum Virtualize

Try it out! From the Storage menu, click Block Storage Systems. Double-click a storage system that runs IBM Spectrum Virtualize, and then click Performance in the General section of the navigation pane.



By default, the charts compare the performance metrics that were collected each hour over the last 24 hours with the guideline values. You can use the calendar, which is next to the title of the page, to change the date and compare current values with historical values.

In most of the charts, a horizontal line is used to indicate the guideline value for the metric. If your devices report prolonged periods of slow response times, you can check whether the performance values for your devices are over the guideline values.

For example, you experience slow send response times for a node in a cluster. You check the chart that tracks the send response times for the nodes and see that one of the node's response times is higher than the guideline value. You can then take remedial action such as balancing the workload of the nodes across the cluster. Alternatively, you can move some of the workloads to other storage systems.

The following key performance indicators are analyzed:

Max Cache Fullness by Pool

The maximum amount of the lower cache that the write cache partitions on the nodes that manage the pool are using for write operations. If the value is 100%, one or more cache partitions on one or more pools is full. The operations that pass through the pools with full cache partitions are queued and I/O response times will increase for the volumes in the affected pools. Available in IBM Spectrum Virtualize 7.3 or later.



The guideline value is 80%.



A critical alert for max write cache fullness is automatically generated when the value is equal to or more than 99%.

Overall Port Bandwidth Percentage by Port

The percentage of the port bandwidth that is used for receive and send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.



The guideline value is 50%.

Compare the guideline value for this metric with the values measured for the switch ports. Because a cluster can have many ports, the chart shows only the eight ports with the highest average bandwidth.



A warning alert is automatically generated when the value for port receive bandwidth or port send bandwidth is equal to or more than 75%. A critical alert is generated when the value for port receive bandwidth or port send bandwidth is equal to or more than 85%.

Port-to-Local Node Send Response Time by Node

The average number of milliseconds to complete a send operation to another node that is in the local cluster. This value represents the external response time of the transfers.



The guideline value is 0.6 ms/op.

Port-to-Remote Node Send Response Time by Node

The average number of milliseconds it takes to complete a send operation to a node in the remote cluster. And, the average number of milliseconds it takes to complete a receive operation from a node in the remote cluster. This value represents the external response time of the transfers.

A guideline value isn't available for this metric because response times for copy-service operations can vary widely.

You can compare the response times to identify discrepancies between the response times for the different nodes.

Read Response Time by I/O Group

The average number of milliseconds to complete a read operation.



The guideline value is 15 ms/op.

Write Response Time by I/O Group

The average number of milliseconds to complete a write operation.



The guideline value is 5 ms/op.

Node Utilization Percentage by Node

The average bandwidth percentages for the ports in the node that are actively used for host and MDisk receive and send operations. The average is weighted based on the port speed and the technology limitations of the node hardware.



The guideline value is 60%.

Port Send Delay Time

The average number of milliseconds of delay that occurs on the port for each send operation. The reason for these delays might be a lack of buffer credits.

A guideline value isn't available for this metric because delay times can vary significantly depending on configuration and usage.

Compare the delay times to identify discrepancies between the ports' delay times and any spikes that might correlate with the time of any reported performance problems.

The Port Send Delay Time is shown instead of the Zero Buffer Credit Percentage by Node chart for some IBM FlashSystem storage systems, such as FlashSystem 9110.

Zero Buffer Credit Percentage by Node

The amount of time, as a percentage, that the port wasn't able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that the node was reset. In Fibre Channel (FC) technology, buffer-to-buffer credit is used to control the flow of frames between ports.

Information about zero buffer credit is only collected and analyzed for 8 Gbps FC ports.



The guideline value is 20%.

Tip: When you add a storage system, a default alert policy is assigned to the storage system. For example, when you add IBM SAN Volume Controller, the default policy for IBM SAN Volume Controller is automatically assigned to the storage system. To find out which alerts are in the default policy, click the Settings menu and click Alert Policies. To see the alerts that are defined for the policy, double-click the default policy.

More actions

Review the following resources to find out:

- [How to investigate the causes of slow drain.](#)
- [Defining alerts.](#)
- [How to define warning and critical thresholds for alerts.](#)

Restrictions

The following charts aren't available for storage systems that run IBM Spectrum Virtualize and use the iSCSI protocol to connect to storage systems:

- Overall Port Bandwidth Percentage by Port
- Node Utilization Percentage by Node
- Port Send Delay Time

Viewing performance alerts

View the alerts that were generated when the measured value of a performance metric falls outside of the specified threshold for a resource.

Before you begin

To generate performance alerts, you must collect performance data and define performance alerts for that resource, or for the alert policy that manages the resource. For information about how to collect performance data, see [Creating performance monitors in IBM Spectrum Control](#). For information about how to define performance alerts, see [Defining alert definitions for performance changes](#).

Procedure

1. Complete one of the following tasks to view performance alerts:

Task	Steps
View the performance alerts for a specific storage system	<ol style="list-style-type: none">a. In the menu bar, choose one of the following options:<ul style="list-style-type: none">• Storage > Block Storage Systems• Storage > File Storage Systemsb. Double-click a storage system.c. Click Alerts in the General section.
View all the performance alerts that were generated for switches	<ol style="list-style-type: none">a. In the menu bar, select Network > Switches.b. Click the Alerts tab.
View the performance alerts for a specific switch	<ol style="list-style-type: none">a. In the menu bar, select Network > Switches.b. Double-click a switch.c. Click Alerts in the General section.

2. Right-click the performance alert that you want to view and select View Alert.


In this view, a chart shows the performance of the resource that violated the threshold and generated the alert. The time range of the chart is 2 hours before and 2 hours after the violation occurred. The lines on the chart represent the following values:

- The blue line represents the performance of the resource for the selected metric.
- The red line represents the threshold that was violated.
- The dotted line represents the time when the violation occurred.

For certain storage systems, a ranked list of the 25 volumes that were the top contributors to the workload on the resource where the threshold violation occurred might be displayed below the chart. This list is shown under the following conditions:

- The performance threshold was violated on an internal resource other than a port.
- The performance threshold was violated on any of the following storage systems:
 - DS8000®
 - XIV®
 - IBM Spectrum Accelerate
 - SAN Volume Controller
 - Storwize® V3500
 - Storwize V3700
 - Storwize V7000
 - Storwize V7000 Unified

You can also view the servers or hosts to which those volumes are assigned.

3. Optional: To access the full performance view of a resource that violated the threshold, click the open performance view icon  on the chart.

Exporting performance data for storage systems and fabrics

To help resolve performance issues with storage systems and fabrics, you can export performance data for the resources to a compressed file. If you contact IBM® Support to help you analyze the problem, you might be asked to send this file.

You can export performance data for the following types of resource: storage systems, nodes, modules, I/O groups, host connections, pools, RAID arrays, managed disks, disks, volumes, file systems, switches, ports, and inter-switch connections.

The summarization type for the exported data can be sample, hourly, or daily.

The performance data is created as CSV files, one file for each resource type. By default, the files are compressed as .zip files.

- [Exporting performance data by using the GUI](#)

You can export performance data for a managed resource. If you contact IBM Support to help you analyze a performance problem with storage systems or fabrics, you might be asked to send this data.

- [Exporting performance data by using a script](#)

Use the exportPerformanceData script to export performance data for managed resources. If you contact IBM Support to help you analyze a performance problem with storage systems or fabrics, you might be asked to send this data.

Exporting performance data by using the GUI

You can export performance data for a managed resource. If you contact IBM® Support to help you analyze a performance problem with storage systems or fabrics, you might be asked to send this data.

Before you begin

The performance data might be large, especially if the data is for storage systems that have many volumes, or the performance monitors are running with a 1-minute sampling frequency. If the time range for the data is greater than 12 hours, volume data and 1-minute sample data is automatically excluded from the performance data even if it is available. To include volume data and 1-minute sample data, select the Advanced package option when you export performance data.

About this task

When you export performance data, you can specify a time range to export performance data for. The time range cannot exceed the history retention limit for sample performance data. By default, this history retention limit is two weeks.

To export hourly or daily performance data, use the exportPerformanceData script. However, the time range still cannot exceed the history retention limits for the type of performance data.

Procedure

1. In the menu bar, select the type of storage system.
For example, to create a compressed file for a block storage system, go to Storage > Block Storage Systems.
To create a compressed file for a fabric, go to Network > Fabrics.
2. Right-click the storage resource, and then click Export Performance data.

Related tasks

- [Configuring history and data retention](#)

Exporting performance data by using a script

Use the exportPerformanceData script to export performance data for managed resources. If you contact IBM® Support to help you analyze a performance problem with storage systems or fabrics, you might be asked to send this data.

Running the exportPerformanceData script

You can run the exportPerformanceData script on the server where IBM Spectrum® Control is installed. The script is run from the *installation_directory/scripts* directory on the server. You can export the performance data for one or more managed resources.

The exportPerformanceData script can be run on Windows and AIX®/Linux® servers:

- Windows default location: `C:\Program Files\IBM\TPC\scripts\exportPerformanceData.bat`
- AIX/Linux default location: `/opt/IBM/TPC/scripts/exportPerformanceData.sh` or `/usr/IBM/TPC/scripts/exportPerformanceData.sh`

The following arguments are required when you run the **exportPerformanceData** script:

```
-user user_name
    Specifies an IBM Spectrum Control user ID.
-pwd password
    Specifies the password for the IBM Spectrum Control user ID.
-resNames resource_name1 resource_name2 resource_name3...
    Specifies the names of storage system, switch, or fabric names as displayed in the UI, separated by spaces. If the resources have spaces in the names, then you must enclose the resource names in double quotation marks ("").
```

To view the optional parameters that the script uses, run the script with the -help option.

Examples

To collect and export performance data to a compressed file for a storage system:

```
>cd c:\Program Files\IBM\TPC\scripts\
>exportPerformanceData -user scUser1 -pwd scPwd1
    -out "c:\web_data\performance_reports\Array1.20160402"
    -resTypes storageSystem,Nodes,Pools,Ports
    -start "2016-02-04 00:00:00" -end "2016-02-04 23:59:59"
```

```
-summType hourly -advPkg yes -zip no  
-resNames "Production Array 1"
```

This command creates a directory Array1.20160402 in c:\web data\performance reports\ if the directory does not exist already. The command writes to that directory the CSV files that contain the hourly performance data for the whole storage system, nodes, pools, and ports. The data is for the storage system named Production Array 1. The time range of the data is from Apr 2, 2016 0:00 AM to Apr 2, 2016 11:59 PM. One minute sample and volume data are included where they are available in the specified time range.

Performance metrics

IBM Spectrum® Control can collect information about the performance of storage systems and switches. This information includes metrics that measure the performance of volumes, ports, and disks. You can view performance information or specify alerts for when the performance of a resource falls outside a specified range.

You can view performance metrics at the following locations:

Resource pages

To access a view of performance metrics for a resource, go to the following pages:

- Storage systems page, Storage system details page
- Volumes page
- Pools page
- Switches page, Switch details page
- Server details page > Related Resources
- Hypervisor details page > Related Resources

Cognos® Analytics reports

To view detailed performance reports for multiple resources in the Cognos Analytics reporting tool, go to the following URL:

http://hostname:port/bi

The *hostname* is the server that Cognos Analytics is running on. The default value for *port* is 9300.

For example, to view performance reports, go to the following URL:

http://myhostname:9300/bi

- [Performance metrics for DS8000](#)
Monitor the performance metrics that are collected for DS8000® storage systems.
- [Performance metrics for resources that run IBM Spectrum Virtualize](#)
Monitor the performance metrics that are collected for IBM Spectrum Virtualize storage systems.
- [Performance metrics for XIV, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R](#)
Monitor the performance metrics that are collected for XIV® systems, IBM Spectrum Accelerate, IBM FlashSystem® A9000, and IBM FlashSystem A9000R.
- [Performance metrics for IBM Spectrum Scale](#)
Monitor the performance metrics that are collected for IBM Spectrum Scale storage systems.
- [Performance metrics for IBM FlashSystem 900](#)
To review trends in performance for IBM FlashSystem 900 storage systems, you add performance metrics to performance charts. Use the charts to monitor the performance of the storage systems.
- [Performance metrics for Dell EMC storage systems](#)
Monitor the performance metrics that are collected for Unity, VMAX, and VNX storage systems.
- [Performance metrics for Hitachi VSP storage systems](#)
Monitor the performance metrics that are collected for storage systems.
- [Performance metrics for NetApp storage systems](#)
Monitor the performance metrics that are collected for NetApp storage systems that are running ONTAP 9.
- [Performance metrics for Pure storage systems](#)
Monitor the performance metrics that are collected for your Pure FlashArray//M and FlashArray//X storage systems.
- [Performance metrics for other storage systems](#)
Monitor the performance of storage systems that are managed by SMI-S providers. You collect performance metadata from storage systems that are compliant with SMI-S 1.5 or later. The SMI-S providers can be referred to by various names, such as CIM agent, CIMOM (CIM Object Manager) agent, or SMI-S agent.
- [Performance metrics for switches](#)
Monitor the performance metrics that are collected for physical switches, switch ports, and inter-switch connections.

Related tasks

- [Collecting performance data by using IBM Spectrum Control performance monitors](#)
- [Configuring the collection of performance data for IBM Spectrum Scale](#)
- [Defining alert definitions for performance changes](#)

Performance metrics for DS8000

Monitor the performance metrics that are collected for DS8000® storage systems.

Overview

Performance metrics are available for the following resources:

- [Volume metrics](#)
- [Disk metrics](#)

- [Port metrics](#)

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips and the table footnotes.

Volume metrics

Volume performance metrics are divided into the following categories:

- [Key volume metrics](#)
- [I/O rate metrics](#)
- [Cache hit percentage metrics](#)
- [Average transfer size and HPF I/O metrics](#)

Tip: Unless otherwise noted, you can view the volume metrics in [Table 1](#), [Table 2](#), [Table 3](#), and [Table 4](#) for the following resources:

- Host connections
- Nodes
- Pools
- Storage systems
- Volumes

Table 1. Key volume metrics

Metric	Description
Cache Holding Time ¹	The average number of seconds that I/O data for a storage system node is held in the cache. A short cache-holding time indicates adverse performance.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Pool Activity Score ²	The activity level of pools, which is set to the following value: $[\text{Read I/O Rate} \times (1 - \text{Read I/O Cache Hit \%})] \div \text{Total Pool Capacity}$
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Volume Utilization ³	The average percentage of time that the volume is busy.
Write-Cache Delay Percentage	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.
Notes:	
1. This metric is only available when you view the performance of nodes or storage systems.	
2. This metric is only available when you view the performance of pools.	
3. This metric is only available when you view the performance of volumes.	

Table 2. I/O Rates

Metric	Description
Average Transfer Rate (Cache-to-Disk) ¹	The average number of sectors or tracks per second that are transferred from the cache to the disks.
Average Transfer Rate (Disk-to-Cache) ¹	The percentage of cache hits for record-mode read operations. For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.
High Performance FICON® (Read)	The average number of read operations per second that are issued by the High Performance FICON feature of the storage system.
High Performance FICON (Write)	The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.
High Performance FICON (Total)	The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.
Normal I/O Rate (Read)	The average number of nonsequential read operations per second.
Normal I/O Rate (Write)	The average number of nonsequential write operations per second.
Normal I/O Rate (Total)	The average number of nonsequential I/O operations per second. This value includes both read and write operations.
PPRC Transfer Rate	The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.
Record Mode Read I/O Rate ¹	The average number of I/O operations per second for record-mode read operations. For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.
Sequential I/O Rate (Read)	The average number of sequential read operations per second.
Sequential I/O Rate (Write)	The average number of sequential write operations per second.
Sequential I/O Rate (Total)	The average number of sequential I/O operations per second. This value includes both read and write operations.
Write-Cache Delay I/O Rate	The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.

Metric	Description
Notes:	
1. This metric is only available when you view the performance of volumes, pools, nodes, and host connections.	

Table 3. Cache hit percentages

Metric	Description
Normal I/O Cache Hits (Read)	The percentage of nonsequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Normal I/O Cache Hits (Write)	The percentage of nonsequential write operations that are handled in the cache.
Normal I/O Cache Hits (Total)	The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write)	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Record Mode Read Cache Hit Percentage ¹	The percentage of cache hits for record-mode read operations. For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.
Sequential I/O Cache Hits (Read)	The percentage of sequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Sequential I/O Cache Hits (Write)	The percentage of sequential write operations that are handled in the cache.
Sequential I/O Cache Hits (Total)	The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Note:	
1. This metric is only available when you view the performance of volumes, pools, nodes, and host connections.	

Table 4. Average transfer size, HPF I/O, and miscellaneous metrics

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.
HPF I/O Percentage	The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

Disk metrics

Disk performance metrics are divided into the following categories:

- [Key disk metrics](#)
- [Transfer size metrics](#)

Tip: Unless otherwise noted, you can view the disk metrics in [Table 5](#) and [Table 6](#) for the following resources:

- Nodes
- Pools
- RAID arrays
- Ranks
- Storage systems

Table 5. Key disk metrics

Metric	Description
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Metric	Description
Disk Utilization Percentage ¹	<p>The average percentage of time that the disks that are associated with an array are busy. No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value N/A is displayed. This limitation applies only to DS8000 storage systems. If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.</p> <p>Tip: Some highly sequential workloads such as batch or backup processing might continually exceed the threshold because they drive the arrays to high utilization percentages. For these types of workloads, a high utilization indicates that the work is being performed very efficiently and is not a cause for concern.</p> <p>For DS8000, this metric is available only for 8.5.0 or later. For earlier versions, the value N/A is shown in the performance table view and the metric is not shown on the performance chart.</p>
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Response Time (Read)	The average number of milliseconds for the back-end storage resources to respond to a read operation.
Response Time (Write)	The average number of milliseconds for the back-end storage resources to respond to a write operation.
Response Time (Overall)	The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.
<p>Note:</p> <p>1. This metric is only available when you view the performance of RAID arrays.</p>	

Table 6. Transfer size metrics

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Average Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics

Port performance metrics are divided into the following categories:

- [Key port metrics](#)
- [Port I/O rate metrics](#)
- [Port data rate metrics](#)
- [Port response time metrics](#)
- [Error rate metrics](#)
- [Remote mirror metrics](#)
- [Transfer size metrics](#)

Tip: Unless otherwise noted, you can view the port metrics in [Table 7](#), [Table 11](#), [Table 12](#), and [Table 13](#) for ports and storage systems. The port metrics in [Table 8](#), [Table 9](#), [Table 10](#) are only available for ports.

Table 7. Key port metrics

Metric	Description
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Send)	The average rate at which data is sent from the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Total)	The average number of send operations and receive operations per second.
Response Time (Receive)	The average number of milliseconds to complete a receive operation.
Response Time (Send)	The average number of milliseconds to complete a send operation.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.
Port Utilization (Receive) ¹	The average percentage of time that the port is busy receiving data.
Port Utilization (Send) ¹	The average percentage of time that the port is busy sending data.
Port Utilization (Overall) ¹	The average percentage of time that the port is busy sending or receiving data.
Bandwidth (Receive) ¹	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Send) ¹	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Overall) ¹	The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Metric	Description
Note:	
1. This metric is only available when you view the performance of ports.	

Table 8. Port I/O rate metrics

Metric	Description
FICON I/O Rate (Send)	The average number of send operations per second for Fibre Channel connection (FICON) usage.
FICON I/O Rate (Receive)	The average number of receive operations per second for FICON usage.
FICON I/O Rate (Total)	The average number of send and receive operations per second for FICON usage.
FCP I/O Rate (Send)	The average number of send operations per second for Fibre Channel Protocol (FCP) usage.
FCP I/O Rate (Receive)	The average number of receive operations per second for FCP usage.
FCP I/O Rate (Total)	The average number of send operations and receive operations per second for FCP usage.

Table 9. Port data rates

Metric	Description
FICON Data Rate (Send)	The average number of MiB per second that is sent for FICON usage.
FICON Data Rate (Receive)	The average number of MiB per second that is received for FICON usage.
FICON Data Rate (Total)	The average number of MiB per second that is transferred for FICON usage. This value includes both send and receive FICON operations.
FCP Data Rate (Send)	The average number of MiB per second that are sent for FCP usage.
FCP Data Rate (Receive)	The average number of MiB per second that are received for FCP usage.
FCP Data Rate (Total)	The average number of MiB per second that are transferred for FCP usage. This value includes both send and receive FCP operations.

Table 10. Port response times

Metric	Description
FICON Response Time (Send)	The average number of milliseconds to complete a send operation for FICON usage.
FICON Response Time (Receive)	The average number of milliseconds to complete a receive operation for FICON usage.
FICON Response Time (Overall)	The average number of milliseconds to complete a send or receive operation for FICON usage. This value includes both send and receive FICON operations.
FCP Response Time (Send)	The average number of milliseconds to complete a send operation for FCP usage.
FCP Response Time (Receive)	The average number of milliseconds to complete a receive operation for FCP usage.
FCP Response Time (Overall)	The average number of milliseconds to complete a send or receive operation for FCP usage.

Table 11. Error rate metrics.

Tip: You can view the following metrics for ports and storage systems.

Metric	Description
Frame Errors (CRC Errors)	The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.
Frame Errors (Error Frame)	The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.
Frame Errors (Invalid Relative Offset Rate)	The average number of times per second that frames are received with an invalid relative offset in the frame header.
Link Errors (Invalid Link Transmission)	The average number of times per second that an invalid transmission word was detected by the port while the link did not experience any signal or synchronization loss.
Link Errors (Invalid Transmission Words)	The average number of bit errors per second that are detected.
Link Errors (Link Failures)	The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.
Link Errors (Primitive Sequence Protocol Error Rate)	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.
Link Errors (Sequence Timeouts)	The average number of times per second that the port detects a timeout condition after the port receives a sequence initiative for a Fibre Channel exchange.
Link Errors (Signal Loss)	The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.
Link Errors (Sync Loss)	The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.
Port Protocol Errors (Credit Recovery Link Resets)	The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

Metric	Description
Port Protocol Errors (Duplicate Frames)	The average number of duplicate frames per second that are received. A duplicate frame is a frame that the system previously processed for the port.
Port Protocol Errors (Link Reset Received)	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.
Port Protocol Errors (Link Reset Transmitted)	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.
Port Protocol Errors (Out of Order ACK)	The average number of times per second that an out-of-order acknowledge (ACK) frame is detected. An ACK frame is used for end-to-end flow control and is sent to verify receipt of a frame.
Port Protocol Errors (Out of Order Data)	The average number of times per second that an out-of-order frame is detected.
Note:	
1. This metric is only available when you view the performance of ports.	

Table 12. Remote mirror metrics

Metric	Description
PPRC Data Rate (Receive)	The average number of MiB per second that are received by using the Peer-to-Peer Remote Copy (PPRC) protocol.
PPRC Data Rate (Send)	The average number of MiB per second that are sent by using the PPRC protocol.
PPRC Data Rate (Total)	The average number of MiB per second that are transferred by using the PPRC protocol. This value includes both send and receive PPRC operations.
PPRC I/O Rate (Receive)	The average number of operations per second that are received by using the PPRC protocol.
PPRC I/O Rate (Send)	The average number of operations per second that are sent by using the PPRC protocol.
PPRC I/O Rate (Total)	The average number of send operations and receive operations per second using the PPRC protocol.
PPRC Response Time (Receive)	The average number of milliseconds to complete a receive operation by using the PPRC protocol.
PPRC Response Time (Send)	The average number of milliseconds to complete a send operation by using the PPRC protocol.
PPRC Response Time (Overall)	The average number of milliseconds to complete a send or receive operation by using the PPRC protocol.

Table 13. Transfer size metrics

Metric	Description
Average Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Average Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for resources that run IBM Spectrum Virtualize

Monitor the performance metrics that are collected for IBM Spectrum Virtualize storage systems.

Overview

In this documentation, IBM Spectrum Virtualize is used to refer collectively to IBM® SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, and IBM Storwize® storage systems, and to IBM FlashSystem® devices that run IBM Spectrum Virtualize.

Definitions are provided for the performance metrics that are collected for the following storage systems:

- FlashSystem 5000
- FlashSystem 5100
- FlashSystem 7200
- FlashSystem 7300
- FlashSystem V9000
- FlashSystem 9100
- FlashSystem 9200
- FlashSystem 9500
- SAN Volume Controller
- IBM Spectrum Virtualize for Public Cloud
- Storwize V3500
- Storwize V3700
- Storwize V5000
- Storwize V7000
- Storwize V7000 Unified (block storage only)

The following terms are used in the performance metrics for these storage systems:

Stage

To write data from a disk to the cache. The data is not prefetched data.

Destage

To write data from the cache to a disk.

Prestage

To write prefetched data from a disk to the cache.

The performance metrics are divided into the following categories:

- [Volume performance metrics](#)
- [Disk performance metrics](#)
- [Pool performance metrics](#)
- [FC Port performance metrics](#)
- [IP Workload performance metrics](#)
- [Node performance metrics](#)

Volume performance metrics

Volume performance metrics are divided into the following categories:

- [Key metrics for volumes](#)
- [I/O rate metrics for volumes](#)
- [Cache hit percentage metrics for volumes](#)
- [Response time metrics for volumes](#)
- [Remote mirror metrics for volumes](#)
- [Volume cache \(VC\) metrics for volumes](#)
- [Volume copy cache \(VCC\) metrics for volumes](#)
- [Compression metrics for volumes](#)
- [Miscellaneous metrics for volumes](#)
- [Legacy cache metrics for volumes](#)

Tip:

Unless otherwise noted, you can view the volume metrics in [Table 1](#), [Table 2](#), [Table 3](#), [Table 4](#), [Table 5](#), [Table 6](#), [Table 7](#), [Table 9](#), and [Table 10](#) for the following resources:

- Host connections
- I/O groups
- Nodes
- Pools
- Storage systems
- Volumes

Table 1. Key metrics for volumes

Metric	Definition
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Unmap) ¹	The average number of MiBs per second that were unmapped. This metric corresponds to the collected ub statistic.
Data Rate (Total)	The average number of MiB per second that are transferred for read, write, and unmap operations. ¹
Overall Host Attributed Response Time Percentage	The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently. For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Unmap) ¹	The average number of unmap operations per second. This metric corresponds to the collected uo statistic.
Overall I/O Rate (Total)	The average number of nonsequential I/O operations per second. This value includes read, write, and unmap operations. ¹
Pool Activity Score ²	The activity level of pools, which is set to the following value: $[\text{Read I/O Rate} \times (1 - \text{Read I/O Cache Hit \%})] \div \text{Total Pool Capacity}$
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Unmap) ¹	The average number of milliseconds required to complete an unmap operation. This metric corresponds to the collected ul statistic.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Metric	Definition
Notes:	
<ol style="list-style-type: none"> 1. This metric applies only to storage systems that are running IBM Spectrum Virtualize 8.1.1 or later. To view details about collected statistics, see Starting statistics collection. 2. This metric is also available when you view the performance of pools. 3. This metric is only available when you view the performance of volumes. 	

Table 2. I/O rate metrics for volumes

Metric	Definition
Transfer Rate (Cache-to-Disk)*	The average number of sectors or tracks per second that are transferred from the cache to the disks.
Transfer Rate (Disk-to-Cache)*	The average number of sectors or tracks per second that are transferred per second from the disks to the cache.
Unaligned Unmap I/O Rate	The average number of volume unmap operations per second that are not aligned on an 8K boundary. This metric corresponds to the collected uou statistic. This metric applies only to storage systems that are running IBM Spectrum Virtualize 8.1.1 or later. To view details about collected statistics, see Starting statistics collection .
Note: *This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.	

Table 3. Cache hit percentages metrics for volumes

Metric	Definition
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write)	The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Table 4. Response time metrics for volumes

Metric	Definition
Peak Response Time (Read)	The worst response time measured for a read operation in the sample interval.
Peak Response Time (Write)	The worst response time measured for a write operation in the sample interval.
Peak Response Time (Unmap)	The worst response time measured for an unmap operation in the sample interval. This metric corresponds to the collected ulw statistic. This metric applies only to storage systems that are running IBM Spectrum Virtualize 8.1.1 or later. To view details about collected statistics, see Starting statistics collection .

Table 5. Remote mirror metrics for volumes

Metric	Definition
Global Mirror (Overlapping Write I/O Rate)	The average number of overlapping write operations per second that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.
Global Mirror (Overlapping Write Percentage)	The percentage of overlapping write operations that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.
Global Mirror (Secondary Write Lag)	The average number of additional milliseconds that it takes to service each secondary write operation for Global Mirror. This value does not include the time to service the primary write operations. You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.
Global Mirror (Write I/O Rate)	The average number of write operations per second that are issued to the Global Mirror secondary site.

Volume cache (VC) metrics are only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is 7.3 or later.

Tip:

The volume cache is sometimes referred to as *upper cache*.

Table 6. Volume cache (VC) metrics for volumes

Metric	Definition
Cache Hits (Dirty Writes)	The percentage of all cache write hits that occur on data that is marked as modified in the volume cache. This value represents how effectively write operations are coalesced before the data is written to disk.
Cache Hits (Read)	The percentage of read operations that find data in the volume cache.
Cache Hits (Write)	The percentage of cache hits for write operations that are handled in the volume cache.
Fast-Write Write Data Rate	The average number of MiB per second that were written to disk in fast-write made in the upper cache. Use this information to help identify the source of back-end overloading, measure the workloads that exit the upper caches, and detect IO amplification in general.
I/O Rate (Destage)	The average number of cache-to-disk transfer operations per second that are processed in the volume cache.
I/O Rate (Read)	The average number of read operations per second that are processed in the volume cache. This value includes operations that are started by hosts or by remote replication sources.
I/O Rate (Write)	The average number of write operations per second that are processed by the volume cache. This value includes operations that are started by hosts or by remote replication sources.
Response Time (Destage)	The average number of milliseconds that it took to complete each destage operation in the volume cache. That is, the time that it took to do write operations from the volume cache to the disk.
Response Time (Stage)	The average number of milliseconds that it took to complete each stage operation in the volume cache. That is, the time that it took to do read operations from the disk to the volume cache.

Metric	Definition
Write Delay Percentage (Flush-through)	The percentage of write operations that are written to disk in flush-through mode in the volume cache.
Write Delay Percentage (Total Delay)	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions in the volume cache. The value is a percentage of all operations.
Write Delay Percentage (Write-through)	The percentage of write operations that are written to disk in write-through mode in the volume cache.
Write Delay Rate (Flush-through)	The average number of tracks per second that are written to disk in flush-through mode in the volume cache.
Write Delay Rate (Total Delay)	The average number of I/O operations per second that are delayed. The delay might occur because of space constraints in the write cache, or because of other conditions in the volume cache. The value is an average of all operations.
Write Delay Rate (Write-through)	The average number of sectors per second that are written to disk in write-through mode in the volume cache.

Volume copy cache (VCC) metrics are only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is 7.3 or later.
Tip: The volume copy cache is sometimes referred to as *lower cache*.

Table 7. Volume copy cache (VCC) metrics for volumes

Metric	Definition
Cache Hits (Dirty Writes)	The percentage of all cache write hits that occur on data that is marked as modified in the volume copy cache. This value represents how effectively write operations are coalesced before the data is written to disk.
Cache Hits (Read-ahead)*	The percentage of all read cache hits that occur on pre-staged data.
Cache Hits (Read)	The percentage of read operations that find data in the volume copy cache.
Cache Hits (Write)	The percentage of cache hits for write operations that are handled in the volume copy cache.
Fast Write Data Rate	The average number of MiB per second that were written to disk in fast-write mode in the lower cache. Use this information to help identify the source of back-end overloading, measure the workloads that exit the lower copy caches, and detect IO amplification in general.
I/O Rate (Destage)	The average number of cache-to-disk transfer operations per second that are processed in the volume copy cache.
I/O Rate (Prestage)	The average number of prefetch disk-to-cache transfer operations per second that are processed in the volume copy cache.
I/O Rate (Read)	The average number of read operations per second that are processed in the volume copy cache. This value includes read operations that are associated with FlashCopy® services, volume mirroring, and other internal processes. This value might also include some operations that are passed from the volume cache.
I/O Rate (Write)	The average number of write operations per second that are processed by the volume copy cache. This value includes read operations that are associated with FlashCopy services, volume mirroring, and other internal processes. This value might also include some operations that are passed from the volume cache.
Response Time (Destage)	The average number of milliseconds that it took to complete each destage operation in the volume copy cache. That is, the time that it took to do write operations from the volume copy cache to the disk.
Response Time (Prestage)	The average number of milliseconds that it took to complete each prestage operation in the volume copy cache. That is, the time that it took to prefetch data from the disk into the volume copy cache.
Response Time (Stage)	The average number of milliseconds that it took to complete each stage operation in the volume copy cache. That is, the time that it took to do read operations from the disk to the volume copy cache.
Transfer Rates (Cache-to-Disk)	The average number of sectors that are transferred per second from the volume copy cache to the disks.
Transfer Rates (Disk-to-Cache)	The average number of sectors that are transferred per second from the disks to the volume copy cache.
Write Delay Percentage (Flush-through)	The percentage of write operations that are written to disk in flush-through mode in the volume copy cache.
Write Delay Percentage (Total Delay)	The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions in the volume copy cache. The value is a percentage of all operations.
Write Delay Percentage (Write-through)	The percentage of write operations that are written to disk in write-through mode in the volume copy cache.
Write Delay Rate (Flush-through)	The average number of sectors per second that are written to disk in flush-through mode in the volume copy cache.
Write Delay Rate (Total Delay)	The average number of I/O operations per second that are delayed. The delay might occur because of space constraints in the write cache, or because of other conditions in the volume copy cache. The value is an average of all operations.
Write Delay Rate (Write-through)	The average number of sectors per second that are written to disk in write-through mode in the volume copy cache.
Note: *This metric is only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is 7.4 or later.	

Note:

Unless otherwise noted, you can view the volume metrics in [Table 8](#) for the following resources:

- Nodes
- I/O groups
- Storage systems

Table 8. Compression metrics for volumes

Metric	Definition
--------	------------

Metric	Definition
Compressed Volumes I/O Rate	The average number of all read and write operations per second for compressed volumes.
Compressed Volumes Data Rate	The average number of MiB per second that were read from or written to compressed volumes.
Compressed Volumes Response Time	The average number of milliseconds to complete an I/O operation for compressed volumes. This value includes both read and write operations.
Uncompressed Volumes I/O Rate	The average number of all read and write operations per second for uncompressed volumes.
Uncompressed Volumes Data Rate	The average number of MiB per second that were read from or written to uncompressed volumes.
Uncompressed Volumes Response Time	The average number of milliseconds to complete an I/O operation for uncompressed volumes. This value includes both read and write operations.

Tip:

Unless otherwise noted, you can view the volume metrics in [Table 9](#) for the following resources:

- Nodes
- I/O groups
- Host connections
- Storage systems

Table 9. Miscellaneous metrics for volumes

Metric	Definition
Cache to Host Transfer Response Time ¹	The average number of milliseconds that is taken to transfer a track from the cache to the host, including any queuing time that occurs because of throttling.
Non-Preferred Node Usage Percentage ²	The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.

Notes:

1. The metric is only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is 7.3 or later.
2. This metric is only available when you view the performance of volumes, I/O groups, and host connections.

Legacy cache metrics are only available for SAN Volume Controller and Storwize block storage systems whose firmware version is earlier than 7.3.

Table 10. Legacy cache metrics for volumes

Metric	Definition
Dirty Write Percentage of Cache Hits	The percentage of all cache write hits that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Disk performance metrics

Disk performance metrics are divided into the following categories:

- [Key metrics for disks](#)
- [Response time metrics for disks](#)
- [Miscellaneous metrics for disks](#)

Unless otherwise noted, you can view disk metrics for the following resources:

- Managed disks
- Pools
- Nodes
- I/O Groups
- Storage systems

Restriction: Performance metadata for managed disks in IBM Spectrum Virtualize for Public Cloud is not yet available.

Table 11. Key metrics for disks

Metric	Definition
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
I/O Rate (Read) ¹	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write) ²	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total) ³	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Response Time (Read)	The average number of milliseconds for the back-end storage resources to respond to a read operation.
Response Time (Write)	The average number of milliseconds for the back-end storage resources to respond to a write operation.
Response Time (Overall)	The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Notes:

1. The performance metrics for I/O Rate (Read) are available for pools, nodes, I/O groups, and storage systems.
2. The performance metrics for I/O Rate (Write) are available for pools, nodes, I/O groups, and storage systems.
3. The performance metrics for I/O Rate (Total) are available for pools, nodes, I/O groups, and storage systems.

Table 12. Response time metrics for disks

Metric	Definition
Queue Time (Read)	The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.
Queue Time (Write)	The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.
Queue Time (Overall)	The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.
Peak Back-end Queue Time (Read)	The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.
Peak Back-end Queue Time (Write)	The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.
Peak Back-end Response Time (Read)	The longest time for a back-end storage resource to respond to a read operation.
Peak Back-end Response Time (Write)	The longest time for a back-end storage resource to respond to a write operation by a node.

Table 13. Miscellaneous metrics for disks

Metric	Definition
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.
Cache Destage (In-Flight I/O)*	The actual number of concurrent operations that are issued from the cache to the managed disk. The number of in-flight operations cannot exceed the target, but also varies over time, depending on the type of I/O operations and on the performance of the managed disk.
Cache Destage (Target I/O)*	The maximum number of concurrent operations that can be issued from the cache to the managed disk. This target changes over time, depending on changes in cache utilization and on changes in the performance of the managed disk.
Note: *This metric is only available when you view the performance of managed disks. The metric is only available for SAN Volume Controller, Storwize, and FlashSystem block storage systems whose firmware version is 7.3 or later.	

Pool performance metrics

Key performance metrics are available for pools.

Unless otherwise noted, you can view pool metrics for the following resources:

- Pools

Table 14. Key metrics for pools

Metric	Definition
Max Write Cache Fullness*	The maximum amount of the lower cache that the write cache partitions on the nodes that manage the pool are using for write operations. If the value is 100%, one or more cache partitions on one or more pools is full. The operations that pass through the pools with full cache partitions will be queued and I/O response times will increase for the volumes in the affected pools.
Write Cache Fullness*	The average amount of the lower cache that the pools' write cache partitions on the nodes are using for write operations. Monitor average cache fullness to identify the pools that are experiencing heavy cache usage.
Note: *This cache fullness metric applies to systems that are running IBM Spectrum Virtualize 7.3 or later.	

FC Port performance metrics

FC Port performance metrics are divided into the following categories:

- [Key metrics for FC ports](#)
- [I/O rate metrics for FC ports](#)
- [Data rate metrics for FC ports](#)
- [Response time metrics for FC ports](#)
- [Error rate metrics for ports](#)
- [Miscellaneous metrics for FC ports](#)

Unless otherwise noted, you can view port metrics for the following resources:

- Ports
- Nodes
- I/O Groups
- Storage systems

Table 15. Key metrics for FC Ports

Metric	Definition
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Send)	The average rate at which data is sent from the port. The rate is measured in MiB per second.

Metric	Definition
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Bandwidth (Receive)	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Send)	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Overall)	The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Table 16. I/O rate metrics for FC Ports

Metric	Definition
Port-to-Disk I/O Rate (Receive)	The average number of exchanges per second that are received from back-end storage resources.
Port-to-Disk I/O Rate (Send)	The average number of IOs per second that are sent from the storage system to the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.
Port-to-Disk I/O Rate (Total)	The average number of IOs per second that are transmitted between the storage system and the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.
Port-to-Host I/O Rate (Receive)	The average number of IOs per second that are received by the storage system from the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.
Port-to-Host I/O Rate (Send)	The average number of IOs per second that are sent by the storage system to the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.
Port-to-Host I/O Rate (Total)	The average number of IOs per second that are transmitted between the storage system and the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.
Port-to-Local Node I/O Rate (Receive)	The average number of IOs per second that are received from other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.
Port-to-Local Node I/O Rate (Send)	The average number of IOs per second that are sent to other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.
Port-to-Local Node I/O Rate (Total)	The average number of IOs per second that are transmitted between the resource and other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.
Port-to-Remote Node I/O Rate (Receive)	The average number of IOs per second that are received from nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.
Port-to-Remote Node I/O Rate (Send)	The average number of IOs per second that are sent to nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.
Port-to-Remote Node I/O Rate (Total)	The average number of IOs per second that are transmitted between the resource and nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Table 17. Data rate metrics for FC Ports

Metric	Definition
Port-to-Disk Data Rate (Receive)	The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.
Port-to-Disk Data Rate (Send)	The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.
Port-to-Disk Data Rate (Total)	The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.
Port-to-Host Data Rate (Receive)	The average rate at which data is received from host computers. The rate is measured in MiB per second.
Port-to-Host Data Rate (Send)	The average rate at which data is sent to host computers. The rate is measured in MiB per second.
Port-to-Host Data Rate (Total)	The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.
Port-to-Local Node Data Rate (Receive)	The average rate at which data is received from other nodes that are in the local cluster. The rate is measured in MiB per second.
Port-to-Local Node Data Rate (Send)	The average rate at which data is sent to other nodes that are in the local cluster. The rate is measured in MiB per second.
Port-to-Local Node Data Rate (Total)	The average rate at which data is transmitted between the component and other nodes that are in the local cluster. The rate is measured in MiB per second.
Port-to-Remote Node Data Rate (Receive)	The average rate at which data is received from nodes that are in the remote cluster. The rate is measured in MiB per second.
Port-to-Remote Node Data Rate (Send)	The average rate at which data is sent to nodes that are in the remote cluster. The rate is measured in MiB per second.
Port-to-Remote Node Data Rate (Total)	The average rate at which data is transmitted between the component and nodes that are in the remote cluster. The rate is measured in MiB per second.

Table 18. Response time metrics for FC Ports

Metric	Definition
Port-to-Local Node Response Time (Receive)	The average number of milliseconds to complete a receive operation from another node that is in the local cluster. This value represents the external response time of the transfers.
Port-to-Local Node Response Time (Send)	The average number of milliseconds to complete a send operation to another node that is in the local cluster. This value represents the external response time of the transfers.
Port-to-Local Node Response Time (Overall)	The average number of milliseconds to complete a send or receive operation with another node that is in the local cluster. This value represents the external response time of the transfers.
Port-to-Remote Node Response Time (Receive)	The average number of milliseconds to complete a receive operation from a node that is in the remote cluster. This value represents the external response time of the transfers.

Metric	Definition
Port-to-Remote Node Response Time (Send)	The average number of milliseconds to complete a send operation to a node that is in the remote cluster. This value represents the external response time of the transfers.
Port-to-Remote Node Response Time (Overall)	The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote cluster. This value represents the external response time of the transfers.

Metrics availability restrictions: The response time metrics are available for nodes, I/O groups, and storage systems.

Table 19. Error rate metrics for FC Ports

Metric	Definition
CRC Error Rate	The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.
Link Errors (Invalid Link Transmission Rate)	The average number of bit errors per second that are detected.
Link Errors (Invalid Transmission Word Rate)	The average number of bit errors per second that are detected.
Link Errors (Link Failures)	The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.
Link Errors (Primitive Sequence Protocol Error Rate)	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.
Link Errors (Signal Loss)	The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.
Link Errors (Sync Loss)	The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.
Port Congestion Index*	The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Port Protocol Errors (Port Send Delay Time)	The average number of milliseconds of delay that occur on the port for each send operation. The reason for these delays might be a lack of buffer credits. You cannot view zero buffer credit performance metrics for 16 Gbps Fibre Channel ports on resources that run IBM Spectrum Virtualize. Use the Port Send Delay Time metric if the Zero Buffer Credit Timer metric is not available.
Port Protocol Errors (Port Send Delay I/O Percentage)	The percentage of send operations where a delay occurred, relative to the total number of send operations that were measured for the port. Use this metric with the Port Send Delay Time metric to distinguish a few long delays from many short delays.
Port Protocol Errors (Zero Buffer Credit Percentage)	The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that metadata was collected. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.
Port Protocol Errors (Zero Buffer Credit Timer)	The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that metadata was collected. If this metric is not available, use the Port Send Delay Time metric instead.
»Total Physical Port Error Rate (cnt/s)«	<p>»The sum of all the physical error rates such as Error Frames, CRC Errors, Short Frames, and Link Failures that are detected on the storage system port.</p> <p>Total Physical Port Error Rate is the sum of the following physical error rates:</p> <ul style="list-style-type: none"> • Error Frame Rate • CRC Error Rate • Short Frame Rate • Long Frame Rate • Bad EOF CRC Error Rate • Link Failure Rate • Loss of Sync Rate • Loss of Signal Rate • Primitive Sequence Protocol Error Rate • Invalid Word Transmission Rate <p>«</p>

Metric	Definition
»Total Logical Port Error Rate (cnt/s)« «	»The sum of all the logical error rates such as F-BSY Frames, F-RJT Frames, Discarded Frames, and Encoding Disparity that are detected on the storage system port. Total Logical Port Error Rate is the sum of the following logical error rates: <ul style="list-style-type: none"> • F-BSY Frame Rate • F-RJT Frame Rate • Discarded Class 3 Frame Rate • Discarded Frame Rate • Link Reset Transmitted Rate • Link Reset Received Rate • Class 3 Send Timeout Frame Rate • Class 3 Receive Timeout Frame Rate • Encoding Disparity «
Note: *The performance metric for Port Congestion Index is only available for ports.	

Table 20. Miscellaneous metrics for FC Ports

Metric	Definition
Port-to-Local Node Queue Time (Receive)	The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local cluster.
Port-to-Local Node Queue Time (Send)	The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local cluster.
Port-to-Local Node Queue Time (Overall)	The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local cluster.
Port-to-Remote Node Queue Time (Receive)	The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote cluster.
Port-to-Remote Node Queue Time (Send)	The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote cluster.
Port-to-Remote Node Queue Time (Overall)	The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote cluster.

IP Workload performance metrics

You can view the following metrics for IP Workload.

Table 21. IP Workload metrics for ports

Metric	Definition
IP Replication Compressed Data Rate (Send)	Average number of mebibytes per second that are transmitted after any compression (if active).
IP Replication Compressed Data Rate (Receive)	Average number of mebibytes per second that are received before any decompression.
IP Replication-to-Remote Node Data Rate (Send)	Average number of mebibytes per second that are transferred to other nodes in other clusters by the IP partnership driver. The rate is measured in MiB/second. Use this metric to measure the transferred data rate between the IP Replication to Remote Node.
IP Replication-to-Remote Node Data Rate (Receive)	Average number of mebibytes per second that are received from other nodes in other clusters by the IP partnership driver. The rate is measured in MiB/second. Use this metric to measure the Received data rate between the IP Replication to Remote Node.
IP Replication-to-Remote Node Data Rate (Total)	Average number of mebibytes per second that are re-transferred to other nodes in other clusters by the IP partnership driver. The rate is measured in MiB/second. Use this metric to measure the total data rate between the IP Replication to Remote Node.
IP Replication Latency	Average round-trip time for the IP partnership link since the last statistics collection period. The rate is measured in milliseconds.
IP Replication Transfer Size (Send)	Average number of mebibytes that are transferred by the IP partnership driver since the last statistics collection period. The rate is measured in MiB. Use this metric to measure the data size that is transferred by IP partnership driver.
IP Replication Transfer Size (Receive)	Average number of mebibytes that are received by the IP partnership driver since the last statistics collection period. The rate is measured in MiB. Use this metric to measure the data size that is received by IP partnership driver.
IP Replication Transfer Size (Total)	Total number of mebibytes that are transferred by the IP partnership driver since the last statistics collection period. The rate is measured in MiB. Use this metric to measure the total data size that is transferred by IP partnership driver.

Node performance metrics

Unless otherwise noted, you can view node metrics for the following resources:

- Nodes
- I/O Groups
- Storage systems

Table 22. Metrics for nodes

Metric	Definition
Compression CPU Utilization (Core 1 to Core 28)	The approximate percentage of time that a processor core was busy with data compression tasks. The performance of each core is shown with a separate metric. Note that the value for this metric will be zero or close to zero if compression accelerator hardware is installed in the nodes.
CPU Utilization (Compression CPU)	The average percentage of time that the processors used for data compression I/O tasks are busy.
CPU Utilization (System CPU)	The average percentage of time that the processors on nodes are busy doing system I/O tasks.

Metric	Definition
Data Movement Rate (MiBs) ¹	The capacity, in MiBs per second, of the valid data in a reclaimed volume extent that garbage collection has moved to a new extent in the data reduction pool on the node. The valid data must be moved so that the whole extent can be freed up or reused to write new data. This metric corresponds to the collected mm statistic.
Data Rewrite Rate (MiBs) ¹	The rate, in MiBs per second, at which data is rewritten when a host overwrites data in data reduction pools on the node. The new version of the host data is written to a different location so that the capacity that was used by the previous version of the host data can be freed up and reclaimed. This metric corresponds to the collected cm statistic. You can view this metric for nodes only.
Extent Collection Rate (cnt/s) ¹	The number of volume extents per second that were processed for garbage collection. The reclaimable capacity in the volume extents is collected so that it can be reused in the data reduction pools on the node. This metric corresponds to the collected ext col statistic. You can view this metric for nodes only.
Logical Data Rate (Sent)	The average number of logical mebibytes per second (both local node and remote node) that are sent to other nodes.
Logical Data Rate (Receive)	The average number of logical mebibytes per second (both local node and remote node) that are received from other nodes.
Logical Data Rate (Total)	The average number of logical mebibytes per second (both local node and remote node) that are transmitted between nodes.
Max Read Cache Fullness (%) ²	The maximum amount of the lower cache which the cache partitions of the pools that are managed by the node are using for read operations. If the maximum value for the cache is 100%, the read cache partition for one or more of the pools is full. The read operations that pass through the node to the affected pools will be queued and the I/O response times will increase for the volumes in the affected pools. This metric corresponds to the collected lower cache rfm statistic.
Max Write Cache Fullness (%) ²	The maximum amount of the lower cache which the cache partitions of the pools that are managed by the node are using for write operations. If the maximum value for the cache is 100%, the write cache partition for one or more of the pools is full. The write operations that pass through the node to the affected pools will be queued and the I/O response times will increase for the volumes in the affected pools. This metric corresponds to the collected lower cache wfm statistic.
New Address Write Rate (MiBs) ¹	The capacity in MiBs per second that was used to write the host's data to unallocated addresses in the data reduction pool on the node. Review this metric to determine which hosts are increasing the amount of capacity that is being written to data reduction pools on a node. This metric corresponds to the collected nm statistic. You can view this metric for nodes only.
Node Utilization by Node	The average of the bandwidth percentages of those ports in the node that are actively used for host and MDisk send and receive operations. The average is weighted by port speed and adjusted according to the technology limitations of the node hardware.
Node-to-Local Node Physical Data Rate (Send)	The average number of mebibytes per second that are sent to other nodes in the local cluster.
Node-to-Local Node Physical Data Rate (Receive)	The average number of mebibytes per second that are received from other nodes in the local cluster.
Node-to-Local Node Physical Data Rate (Total)	The average number of mebibytes per second that are transmitted between nodes in the local cluster.
Node-to-Local Node Logical Data Rate (Send)	The average number of logical mebibytes per second that are sent to the other nodes in the local cluster.
Node-to-Local Node Logical Data Rate (Receive)	The average number of logical mebibytes per second that are received from the other nodes in the local cluster.
Node-to-Local Node Logical Data Rate (Total)	The average number of logical mebibytes per second that are transmitted between nodes in the local cluster.
Node-to-Remote Node Logical Data Rate (Send)	The average number of logical mebibytes per second that are sent to nodes in the remote cluster.
Node-to-Remote Node Logical Data Rate (Receive)	The average number of logical mebibytes per second that are received from nodes in the remote cluster.
Node-to-Remote Node Logical Data Rate (Total)	The average number of logical mebibytes per second that are transmitted between nodes in the remote cluster.
Node-to-Remote Node Physical Data Rate (Send)	The average number of mebibytes per second that are sent to nodes in the remote cluster.
Node-to-Remote Node Physical Data Rate (Receive)	The average number of mebibytes per second that are received from nodes in the remote cluster.

Metric	Definition
Node-to-Remote Node Physical Data Rate (Total)	The average number of mebibytes per second that are transmitted between nodes in the remote cluster.
Physical Data Rate (Send)	The average number of physical mebibytes per second that are sent to other nodes.
Physical Data Rate (Receive)	The average number of physical mebibytes per second that are received from other nodes.
Physical Data Rate (Total)	The average number of Physical mebibytes that are transferred per second.
Read Cache Fullness (%) ²	The average amount of the lower cache which the cache partitions of the pools that are managed by the node are using for read operations. Monitor the average cache fullness for read operations to identify the nodes that are experiencing heavy cache usage. This metric corresponds to the collected lower cache rfav statistic.
Reclaimable Capacity (MiB) ¹	The capacity that can be reclaimed in the data reduction pools on the node. This metric corresponds to the collected rec statistic. You can view this metric for nodes only.
Recovered Capacity Rate (MiBs) ¹	The capacity in number of MiBs per second that was recovered by garbage collection for reuse in the data reduction pools on the node. This metric corresponds to the collected rm statistic. You can view this metric for nodes only.
System CPU Utilization » (Core 1 to Core 48)«	The approximate percentage of time that a processor core was busy with system I/O tasks. The performance of each core is shown with a separate metric.
Write Cache Fullness (%) ²	The average amount of the lower cache which the cache partitions of the pools that are managed by the node are using for write operations. Monitor the average cache fullness for write operations to identify the nodes that are experiencing heavy cache usage. This metric corresponds to the collected lower cache wfav statistic.
Notes: <ol style="list-style-type: none"> 1. This garbage collection metric applies to systems that are running IBM Spectrum Virtualize 8.1.2 or later. To view details about collected statistics, see Starting statistics collection. 2. This cache fullness metric applies to systems that are running IBM Spectrum Virtualize 7.3 or later. To view details about collected statistics, see Starting statistics collection. 	

Performance metrics for XIV, IBM Spectrum Accelerate, IBM® FlashSystem A9000, and IBM FlashSystem A9000R

Monitor the performance metrics that are collected for XIV® systems, IBM Spectrum Accelerate, IBM FlashSystem® A9000, and IBM FlashSystem A9000R.

For XIV systems, IBM FlashSystem A9000, and IBM FlashSystem A9000R, performance metrics are available for the following resources:

- [Volume metrics](#)
- [Port metrics](#)

For IBM Spectrum Accelerate, performance metrics are available for the following resource:

- [Volume metrics](#)

Restriction: XIV systems, IBM Spectrum Accelerate, IBM FlashSystem A9000, and IBM FlashSystem A9000R do not track performance statistics for volumes that were never used. Because there are no performance statistics, performance metrics are not shown for these volumes and their related components.

Volume metrics

Volume performance metrics are organized into the following categories:

- [Key storage metrics](#)
- [Cache hit percentage metrics](#)
- [Response time metrics](#)
- [Miscellaneous metrics](#)

Unless otherwise noted, you can view volume metrics for the following resources:

- Volumes
- Pools
- Modules
- Host connections
- Storage systems

Table 1. Key volume metrics

Metric	Description
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Metric	Description
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Important: For XIV systems, all the key storage metrics are available for 10.2.2 or later.	

Table 2. Cache hit percentage metrics

Metric	Description
Data Cache Hits (Read) ²	The percentage of all read data that is read from the cache.
Data Cache Hits (Write) ²	The percentage of all write data that is written to cache slots that are marked as modified.
Data Cache Hits (Overall) ²	The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.
Overall I/O Cache Hits (Read) ¹	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write) ¹	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.
Overall I/O Cache Hits (Total) ¹	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
SSD Read Cache Hits (I/O) ²	The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric. This metric is only available for XIV systems that use solid-state drives as drives for caching.
SSD Read Cache Hits (Data) ²	The percentage of all read data that was read from cache memory on a solid-state drive. This metric is only available for XIV systems that use solid-state drives as drives for caching.
Notes:	
1. For XIV systems, this metric is available in 10.2.2 or later.	
2. For XIV systems, this metric is available in 10.2.4 or later.	

Table 3. Response time metrics

Metric	Description
Cache Hit Response Time (Read)	The average number of milliseconds to complete a read-cache hit operation.
Cache Hit Response Time (Write)	The average number of milliseconds to complete a write-cache hit operation.
Cache Hit Response Time (Overall)	The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.
Cache Miss Response Time (Read)	The average number of milliseconds to complete a read-cache miss operation.
Cache Miss Response Time (Write)	The average number of milliseconds to complete a write-cache miss operation.
Cache Miss Response Time (Overall)	The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.
Response Time by Transfer Size (Small)	The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.
Response Time by Transfer Size (Medium)	The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.
Response Time by Transfer Size (Large)	The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.
Response Time by Transfer Size (Very Large)	The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.
SSD Read Cache Hit Response Time	The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric. This metric is only available for XIV systems that use solid-state drives as drives for caching.
Note: For XIV systems, all the response time metrics are available for 10.2.2 or later.	

Table 4. Miscellaneous metrics

Metric	Description
Average Transfer Size (Read) ¹	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write) ¹	The average number of KiB that are transferred per write operation.

Metric	Description
Average Transfer Size (Overall) ¹	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.
I/O Transfer Size (Small) ²	The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.
I/O Transfer Size (Medium) ²	The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.
I/O Transfer Size (Large) ²	The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.
I/O Transfer Size (Very Large) ²	The percentage of I/O operations with a data transfer size that is greater than 512 KiB.
Data Transfer Size (Small) ²	The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.
Data Transfer Size (Medium) ²	The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.
Data Transfer Size (Large) ²	The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.
Data Transfer Size (Very Large) ²	The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.
Volume Utilization ^{1, 3}	The average percentage of time that the volume is busy.
Pool Activity Score ^{1, 3}	The activity level of pools, which is set to the following value: $[\text{Read I/O Rate} \times (1 - \text{Read I/O Cache Hit \%})] \div \text{Total Pool Capacity}$ The value of this metric is the same for all pools in the storage system.
Notes:	
1. For XIV systems, this metric is available in 10.2.2 or later. 2. For XIV systems, this metric is available in 10.2.4 or later. 3. This metric is only available when you view the performance of pools.	

Port metrics

The following performance metrics for ports are available for XIV systems that are 10.2.4 or later:

- [I/O rate metrics](#)
- [Data rate metrics](#)
- [Response time metrics](#)
- [Miscellaneous port metrics](#)

Restriction: Port metrics are not available for IBM Spectrum Accelerate.

Table 5. I/O rate metrics

Metric	Description
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Total)	The average number of send operations and receive operations per second.

Table 6. Data rate metrics

Metric	Description
Data Rate (Send)	The average rate at which data is sent from the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Table 7. Response time metrics

Metric	Description
Response Time (Send)	The average number of milliseconds to complete a send operation.
Response Time (Receive)	The average number of milliseconds to complete a receive operation.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.

Table 8. Miscellaneous port metrics

Metric	Description
Bandwidth (Send)	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Receive)	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Overall)	The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Performance metrics for IBM Spectrum Scale

Monitor the performance metrics that are collected for IBM Spectrum Scale storage systems.

Overview

The performance metrics are divided into the following categories:

- [Node performance metrics](#)
- [File system performance metrics](#)

Node performance metrics

You can view the following metrics for each IBM Spectrum Scale cluster node.

Table 1. Metrics for nodes

Metric	Definition
CPU Utilization (User)	The average percentage of time that the processors on a node are busy doing user tasks.
CPU Utilization (System)	The average percentage of time that the processors on a node are busy doing system tasks.
CPU Utilization (Total)	The average percentage of time that the processors on a node are busy doing user tasks and system tasks.
Memory Used (Cache and Buffer)	The average percentage of memory on a node that is used for cache and buffer memory.
Memory Used	The average percentage of memory that is used on a node. This value does not include the memory that is used for cache and buffer memory.
Memory Used (Total)	The average percentage of memory that is used on a node. This value includes the memory that is used for cache and buffer memory.
I/O Rate (Read)	The average number of read operations per second.
I/O Rate (Write)	The average number of write operations per second.
I/O Rate (Total)	The average number of read operations and write operations per second.

File system performance metrics

You can view file system metrics for the following resources:

- File systems
 - IBM Spectrum Scale storage systems
- Tip: For a storage system, the metrics contain summary values for all the file systems on the storage system.

Table 2. Metrics for file systems

Metric	Definition
Data Rate (Read)	The average number of MiB per second that are transferred for read operations.
Data Rate (Write)	The average number of MiB per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Maximum Data Rate (Read)	The maximum number of MiB per second that are transferred for read operations.
Maximum Data Rate (Write)	The maximum number of MiB per second that are transferred for write operations.
Maximum Data Rate (Total)	The maximum number of MiB per second that are transferred for read operations and write operations.
I/O Rate (Read)	The average number of read operations per second.
I/O Rate (Write)	The average number of write operations per second.
I/O Rate (Total)	The average number of read operations and write operations per second.
Maximum I/O Rate (Read)	The maximum number of read operations per second.
Maximum I/O Rate (Write)	The maximum number of write operations per second.
Maximum I/O Rate (Total)	The maximum number of read operations and write operations per second.
Response Time (Read)	The average number of milliseconds for the back-end storage resources to respond to a read operation from the file system.
Response Time (Write)	The average number of milliseconds for the back-end storage resources to respond to a write operation from the file system.
Response Time (Overall)	The average number of milliseconds for the back-end storage resources to respond to a read operation or a write operation from the file system.

Performance metrics for IBM FlashSystem 900

To review trends in performance for IBM FlashSystem® 900 storage systems, you add performance metrics to performance charts. Use the charts to monitor the performance of the storage systems.

Overview

The performance metrics are described in the following sections:

- [Volume metrics](#)
- [Disk metrics](#)
- [Port metrics](#)

Restrictions:

- IBM FlashSystem 900 storage systems do not track performance for individual volumes. The volume metrics that are available are provided to allow comparisons across different storage systems, but represent performance that is measured for ports rather than for volumes.
- IBM Spectrum® Control supports only Fibre Channel adapters and ports that are installed on storage systems. For IBM FlashSystem 900 storage systems that have other types of ports, performance data is only available for the drives.
- Response time metrics are not available for all IBM FlashSystem 900 storage systems, depending on the firmware release of the storage system. To collect and view response time metrics, ensure that you have firmware level 1.4.5 or later. For details of the firmware that is supported by IBM Spectrum Control, check the storage information in the [IBM Spectrum Control interoperability matrix](#).

Volume metrics

IBM FlashSystem 900 storage systems do not track the performance of individual volumes. Therefore, the performance metrics that you can view represent the sum of the measured performance of all the ports in the storage system.

Table 1. Metrics for volumes

Metric	Description
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Total)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Disk metrics

You can view the following performance metrics for drives or flash modules on IBM FlashSystem 900 storage systems.

Table 2. Metrics for disks

Metric	Description
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
Flash Health Percentage	The overall flash module health. The health percentage is calculated based on the number of unusable blocks on the flash module.

Port metrics

You can view the following port performance metrics for Fibre Channel ports on IBM FlashSystem 900 storage systems. No performance data is available for InfiniBand or FCoE ports.

Table 3. Metrics for ports

Metric	Description
Bandwidth (Send)	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Receive)	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth (Total)	The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Data Rate (Send)	The average rate at which data is sent from the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Total)	The average number of send operations and receive operations per second.
Response Time (Send)	The average number of milliseconds to complete a send operation.
Response Time (Receive)	The average number of milliseconds to complete a receive operation.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.
Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for Dell EMC storage systems

Monitor the performance metrics that are collected for Unity, VMAX, and VNX storage systems.

Overview

The performance metrics are described in the following sections:

- [Unity storage systems](#)
- [VMAX storage systems](#)
- [VNX storage systems](#)

Restriction: For Dell EMC, performance metrics are only available for block storage systems.

Unity storage systems

The following performance metrics are available for Unity resources:

- [Volume metrics for Unity](#)
- [Disk metrics for Unity](#)
- [Port metrics for Unity](#)
- [Node metrics for Unity](#)
- [File system metrics for Unity](#)

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips and the table footnotes.

Volume metrics for Unity

Volume performance metrics are divided into the following categories:

- [Key volume metrics](#)
- [Cache hit metrics](#)

Tip: You can view the volume metrics in [Key volume metrics](#) for the following resources:

- Volumes
- Pools
- Nodes
- Storage systems

Table 1. Key volume metrics

Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Tip: You can view the volume metrics in [Cache hit metrics](#) for the following resources:

- Nodes
- Storage systems

Table 2. Cache hit metrics

Metric	Description
Overall I/O Cache Hits (Read) %	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write) %	The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.
Overall I/O Cache Hits (Total) %	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk metrics for Unity

You can view the metrics in [Key disk metrics](#) for the following resources:

- Pools
- Disks
- Nodes
- Storage systems

Table 3. Key disk metrics

Metric	Definition
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
Response Time (Overall)	The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics for Unity

You can view the metrics in [Key port metrics](#) for ports, nodes, and storage systems.

Table 4. Key port metrics

Metric	Description
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Send)	The average rate at which data is sent from the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Node metrics for Unity

You can view the metrics in [Key node metrics](#) for nodes and storage systems.

Table 5. Key node metrics

Metric	Description
System CPU	The average percentage of time that the processors on nodes are busy doing system I/O tasks.

File system metrics for Unity

Unless otherwise noted, you can view the performance metrics in [Key file system metrics](#) for the following resources:

- Pools
- Nodes
- File systems
- Storage systems

Tip: For a storage system, the metrics contain summary values for all of the file systems on the storage system.

Table 6. Key file systems metrics

Metric	Definition
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

VMAX storage systems

The following performance metrics are available for VMAX resources:

- [Volume metrics for VMAX](#)
- [Disk metrics for VMAX](#)
- [Port metrics for VMAX](#)

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips and the table footnotes.

Volume metrics for VMAX

Volume performance metrics are divided into the following categories:

- [Key volume metrics](#)
- [Cache hit metrics](#)
- [Transfer size metrics](#)

Tip: Unless otherwise noted, you can view the volume metrics in [Key volume metrics](#), [Cache hit metrics](#), and [Transfer size metrics](#) for the following resources:

- Volumes
- Directors
- Storage systems

Table 7. Key volume metrics

Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read) ¹	The average number of milliseconds to complete a read operation.
Response Time (Write) ¹	The average number of milliseconds to complete a write operation.
Response Time (Overall) ¹	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Volume Utilization ²	The average percentage of time that the volume is busy.
Normal I/O Rate (Read)	The average number of nonsequential read operations per second.
Normal I/O Rate (Write)	The average number of nonsequential write operations per second.
Normal I/O Rate (Total)	The average number of nonsequential I/O operations per second. This value includes both read and write operations.
Sequential I/O Rate (Read)	The average number of sequential read operations per second.
Sequential I/O Rate (Write)	The average number of sequential write operations per second.
Sequential I/O Rate (Total)	The average number of sequential I/O operations per second. This value includes both read and write operations.
Note:	
1. This metric is only available when you view the performance of volumes or storage systems.	
2. This metric is only available when you view the performance of volumes.	

Table 8. Cache hit metrics

Metric	Description
Normal I/O Cache Hits (Read)	The percentage of nonsequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Normal I/O Cache Hits (Write)	The percentage of nonsequential write operations that are handled in the cache.
Normal I/O Cache Hits (Total)	The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Sequential I/O Cache Hits (Read)	The percentage of sequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Sequential I/O Cache Hits (Write)	The percentage of sequential write operations that are handled in the cache.
Sequential I/O Cache Hits (Total)	The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write)	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.

Metric	Description
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 9. Transfer size metrics

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Disk metrics for VMAX

Disk performance metrics for VMAX block storage systems are divided into the following categories:

- [Key disk metrics](#)
- [Transfer size metrics](#)

Table 10. Key disk metrics

Metric	Definition
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
Response Time (Overall)	The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.
Disk Utilization Percentage	The approximate utilization percentage of a particular array over a particular time interval, that is, the average percent of time that the disk associated with the array was busy.

Table 11. Transfer size metrics

Metric	Description
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics for VMAX

You can view the metrics below for ports and storage systems.

Table 12. Port metrics

Metric	Description
I/O Rate (Total)	The average number of send operations and receive operations per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Response Time (Overall)	The average number of milliseconds to complete a send or receive operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

VNX storage systems

The following performance metrics are available for VNX resources:

- [Volume metrics for VNX](#)
- [Disk metrics for VNX](#)
- [Port metrics for VNX](#)

Volume metrics for VNX

Volume performance metrics are divided into the following categories:

- [Key volume metrics](#)
- [Transfer size metrics](#)

Tip: Unless otherwise noted, you can view the volume metrics in [Key volume metrics](#) and [Transfer size metrics](#) for the following resources:

- Volumes
- Nodes
- Storage systems

Table 13. Key volume metrics

Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.

Table 14. Transfer size metrics

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Disk metrics for VNX

Disk performance metrics for VNX storage systems are divided into the following categories:

- [Key disk metrics](#)
- [Transfer size metrics](#)

Table 15. Key disk metrics

Metric	Definition
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Table 16. Transfer size metrics

Metric	Description
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics for VNX

You can view the metrics below for ports and storage systems.

Table 17. Port metrics

Metric	Description
I/O Rate (Total)	The average number of send operations and receive operations per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for Hitachi VSP storage systems

Monitor the performance metrics that are collected for storage systems.

The following performance metrics are available for Hitachi VSP storage systems.

- [Volume metrics](#)
- [Port metrics](#)
- [Node metrics](#)

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips.

Volume metrics

Volume performance metrics are divided into the following categories:

- [Key volume metrics](#)
- [Cache hit metrics](#)

Tip: You can view the volume metrics in [Key volume metrics](#) for the following resources:

- Volumes
- Pools
- Storage systems

Table 1. Key volume metrics

Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Tip: You can view the volume metrics in [Cache hit metrics](#) for the following resources:

- Volumes
- Pools
- Storage systems

Table 2. Cache hit metrics

Metric	Description
Overall I/O Cache Hits (Read) %	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write) %	The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Port metrics

You can view the metrics in [Key port metrics](#) for ports and storage systems.

Table 3. Key port metrics

Metric	Description
I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Response Time (Overall)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Node metrics

You can view the metrics in [Key node metrics](#) for storage systems.

Table 4. Key node metrics

Metric	Description
System CPU	The average percentage of time that the processors on nodes are busy doing system I/O tasks.

Performance metrics for NetApp storage systems

Monitor the performance metrics that are collected for NetApp storage systems that are running ONTAP 9.

The following performance metrics are available for storage systems that are running ONTAP 9:

- [Volume metrics](#)
- [Disk metrics](#)
- [Port metrics](#)
- [Node metrics](#)
- [File system metrics](#)

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips.

Volume metrics

Volume performance metrics are divided into the following categories:

- [Key volume metrics](#)

- [Cache hit metrics](#)

Tip: You can view the volume metrics in [Key volume metrics](#) for the following resources:

- Volumes
- Pools
- Nodes
- Storage systems

Table 1. Key volume metrics

Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Volume cache hit metrics are only available when you view the performance of nodes or storage systems. To check which resources a volume cache hit metric applies to, see the table footnotes.

Table 2. Cache hit metrics

Metric	Description
Overall I/O Cache Hits (Read) % ¹	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Notes:	
1. This metric is only available when you view the performance of storage systems.	

Disk metrics

Disk performance metrics for NetApp storage systems that are running ONTAP 9 are described in [Key disk metrics](#).

Tip: You can view the disk metrics in [Key disk metrics](#) for the following resources:

- Pools
- Disks
- Nodes
- Storage systems

Table 3. Key disk metrics

Metric	Definition
I/O Rate (Read)	The average number of read operations per second that are issued to the back-end storage resources.
I/O Rate (Write)	The average number of write operations per second that are issued to the back-end storage resources.
I/O Rate (Total)	The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.
Data Rate (Read)	The average number of MiB per second that are read from the back-end storage resources.
Data Rate (Write)	The average number of MiB per second that are written to the back-end storage resources.
Data Rate (Total)	The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.
Response Time (Read)	The average number of milliseconds for the back-end storage resources to respond to a read operation.
Response Time (Write)	The average number of milliseconds for the back-end storage resources to respond to a write operation.
Response Time (Overall)	The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.
Transfer Size (Read)	The average number of KiB that are transferred per read operation from the back-end storage resources.
Transfer Size (Write)	The average number of KiB that are transferred per write operation to the back-end storage resources.
Transfer Size (Overall)	The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Port metrics

Port performance metrics are described in [Key port metrics](#).

Tip: You can view the port metrics in [Key port metrics](#) for the following resources:

- Ports
- Nodes
- Storage systems

Table 4. Key port metrics

Metric	Description
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Total)	The average number of send operations and receive operations per second.
Response Time (Send)	The average number of milliseconds to complete a send operation.
Response Time (Receive)	The average number of milliseconds to complete a receive operation.
Response Time (Overall)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Data Rate (Send)	The average rate at which data is sent from the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.
Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Node metrics

You can view the metrics in [Key node metrics](#) for storage systems.

Table 5. Key node metrics

Metric	Description
System CPU	The average percentage of time that the processors on nodes are busy doing system I/O tasks.

File system metrics

Unless otherwise noted, you can view the performance metrics in [Key file system metrics](#) for the following resources:

- Pools
- Nodes
- File systems
- Storage systems

Tip: For a storage system, the metrics contain summary values for all of the file systems on the storage system.

Table 6. Key file systems metrics

Metric	Definition
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Performance metrics for Pure storage systems

Monitor the performance metrics that are collected for your Pure FlashArray//M and FlashArray//X storage systems.

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips.

Volume metrics

Volume performance metrics are described in the [Key volume metrics](#) table.

Tip: You can view the volume metrics in [Key volume metrics](#) for the following resources:

- Volumes
- Storage systems

Table 1. Key volume metrics

Metric	Description
--------	-------------

Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Performance metrics for other storage systems

Monitor the performance of storage systems that are managed by SMI-S providers. You collect performance metadata from storage systems that are compliant with SMI-S 1.5 or later. The SMI-S providers can be referred to by various names, such as CIM agent, CIMOM (CIM Object Manager) agent, or SMI-S agent.

Overview

Performance metrics are available for the following resources:

- [Volume metrics](#)
- [Port metrics](#)

A performance metric might apply to one or more storage resources. To check which resources the performance metric applies to, see the tips and the table footnotes.

Restriction: For NetApp storage systems managed by SMI-S providers, performance monitoring is supported only for volumes and front-end ports. It is not supported at the storage system level.

Volume metrics

Volume performance metrics are divided into the following categories:

- [Key volume metrics](#)
- [Cache hit percentage metrics](#)
- [Average transfer size metrics](#)

Tip: Unless otherwise noted, you can view the volume metrics in [Key volume metrics](#), [Cache hit metrics](#), and [Average transfer size metrics](#) for the following resources:

- Volumes
- Nodes
- Storage systems

Table 1. Key volume metrics

Metric	Description
Overall I/O Rate (Read)	The average number of read operations per second. This value includes both sequential and nonsequential read operations.
Overall I/O Rate (Write)	The average number of write operations per second. This value includes both sequential and nonsequential write operations.
Overall I/O Rate (Total)	The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
Data Rate (Read)	The average number of MiBs per second that are transferred for read operations.
Data Rate (Write)	The average number of MiBs per second that are transferred for write operations.
Data Rate (Total)	The average number of MiB per second that are transferred for read operations and write operations.
Response Time (Read)	The average number of milliseconds to complete a read operation.
Response Time (Write)	The average number of milliseconds to complete a write operation.
Response Time (Overall)	The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
Volume Utilization ¹	The average percentage of time that the volume is busy.
Note:	
1. This metric is only available when you view the performance of volumes.	

Table 2. Cache hit percentages

Metric	Description
Overall I/O Cache Hits (Read)	The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.
Overall I/O Cache Hits (Write)	The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.
Overall I/O Cache Hits (Total)	The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Table 3. Average transfer size metrics

Metric	Description
--------	-------------

Metric	Description
Average Transfer Size (Read)	The average number of KiB that are transferred per read operation.
Average Transfer Size (Write)	The average number of KiB that are transferred per write operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Port metrics

Port performance metrics are divided into the following categories:

- [Key port metrics](#)
- [Transfer size metrics](#)

Tip: You can view the port metrics in [Key port metrics](#) and [Transfer size metrics](#) for ports and storage systems.

Table 4. Key port metrics

Metric	Description
I/O Rate (Receive)	The average number of I/O operations per second for operations in which the port receives data.
I/O Rate (Send)	The average number of I/O operations per second for operations in which data is sent from a port.
I/O Rate (Total)	The average number of send operations and receive operations per second.
Data Rate (Receive)	The average rate at which data is received by the port. The rate is measured in MiB per second.
Data Rate (Send)	The average rate at which data is sent from the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Table 5. Transfer size metrics

Metric	Description
Average Transfer Size (Receive)	The average number of KiB that are transferred per receive operation.
Average Transfer Size (Send)	The average number of KiB that are transferred per send operation.
Average Transfer Size (Overall)	The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Performance metrics for switches

Monitor the performance metrics that are collected for physical switches, switch ports, and inter-switch connections.

Performance metrics for switches are divided into the following categories:

- [Key port metrics](#)
- [I/O rate metrics](#)
- [Peak data rate metrics](#)
- [Frame error rate metrics](#)
- [Port protocol error rate metrics](#)
- [Link error rate metrics](#)
- [Miscellaneous metrics](#)

Attention: To enhance troubleshooting experience for switches, two new switch port metrics namely Total Physical Port Error Rate and Total Logical Port Error Rate are available from 2Q22 update of IBM® Storage Insights. With these newly added metrics, you can quickly identify specific switch port that has high rate of physical and/or logical error rates. The two new metrics are listed in the [Table 1](#).

Table 1. Key port metrics

Metric	Description
Bandwidth Percentage (Send)	The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth Percentage (Receive)	The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Bandwidth Percentage (Overall)	The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.
Data Rate (Send)	The average rate at which data is sent by the port. A send operation is a read operation that is processed, or a write operation that is initiated by the port. The rate is measured in MiB per second.
Data Rate (Receive)	The average rate at which data is received by the port. A receive operation is a write operation that is processed, or a read operation that is initiated by the port. The rate is measured in MiB per second.
Data Rate (Total)	The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Metric	Description
Total Physical Port Error Rate (cnt/s)	<p>The sum of all the physical error rates such as Error Frames, CRC Errors, Short Frames, and Link Failures that are detected on the switch port. Total Physical Port Error Rate is the sum of the following physical error rates:</p> <ul style="list-style-type: none"> Error Frame Rate CRC Error Rate Short Frame Rate Long Frame Rate Bad EOF CRC Error Rate Link Failure Rate Loss of Sync Rate Loss of Signal Rate Primitive Sequence Protocol Error Rate Invalid Word Transmission Rate
Total Logical Port Error Rate (cnt/s)	<p>The sum of all the logical error rates such as F-BSY Frames, F-RJT Frames, Discarded Frames, and Encoding Disparity that are detected on the switch port. Total Logical Port Error Rate is the sum of the following logical error rates:</p> <ul style="list-style-type: none"> F-BSY Frame Rate F-RJT Frame Rate Discarded Class 3 Frame Rate Discarded Frame Rate Link Reset Transmitted Rate Link Reset Received Rate Class 3 Send Timeout Frame Rate Class 3 Receive Timeout Frame Rate Encoding Disparity
Total Port Error Rate	The average number of times per second that an error was detected on the port. This rate is a summation of all the other error rates for the port.

Table 2. I/O rates

Metric	Description
Port Frame Rate (Send)	The average number of frames per second that are sent by the port.
Port Frame Rate (Receive)	The average number of frames per second that are received by the port.
Port Frame Rate (Total)	The average number of frames per second that are transferred. This value includes frames that are sent and received by the port.

Table 3. Peak data rates

Metric	Description
Peak Data Rate (Send) ¹	The highest rate at which data is sent by the port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.
Peak Data Rate (Receive) ¹	The highest rate at which data is received by the port. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Notes:

1. This metric is only available for ports on Brocade switches. You must use an SMI agent to collect data for this metric. In the Storage Management Initiative Specification (SMI-S), peak data rates for a port are reported in the CIM_FcPortRateStatistics class. For SMI-S 1.1, this is an optional class, meaning that it might not be supported by a given vendor. This can result in IBM Spectrum® Control displaying zeros for the peak rates, even when there is non-zero traffic flowing through the port. When a vendor does support this class, the peak data rate represents the peak value of the associated metric counter (as reported by the CIM agent used to access the device) over a recent, vendor-defined measurement window (for example, 20 milliseconds) of the IBM Spectrum Control measurement window (for example, 15 minutes).

Table 4. Frame error rates

Metric	Description
Bad EOF CRC Error Rate ¹	The percentage of nonsequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
CRC Error Rate	The percentage of nonsequential write operations that are handled in the cache.
Discarded Class 3 Frame Rate	The average number of class 3 frames per second that are discarded.
Error Frame Rate ¹	The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.
F-BSY Frame Rate ²	The average number of F-BSY frames per second that are generated. An F-BSY frame is issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N_port is busy.
F-RJT Frame Rate ²	The average number of F-RJT frames per second that are generated. An F-RJT frame is issued by the fabric to indicate that delivery of a frame was denied.
Long Frame Rate	The average number of frames that are received per second that are longer than 2140 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 2140 octet limit is calculated based on the assumption that a frame has 24 header bytes, 4 CRC bytes, and 2112 data bytes.
Short Frame Rate ²	The average number of frames that are received per second that are shorter than 28 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 28 octet limit is calculated based on the assumption that a frame has 24 header bytes, and 4 CRC bytes.

Notes:

1. This metric is only available for ports on Brocade switches.
2. This metric is only available for ports on Cisco switches.

Table 5. Port protocol error rates

Metric	Description
--------	-------------

Metric	Description
Class 3 Receive Timeout Frame Rate ¹	The average number of class 3 frames per second that were discarded after reception because of a timeout condition. The timeout condition occurs while a transmitting port waits for buffer credit from a port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Class 3 Send Timeout Frame Rate ¹	The average number of class 3 frames per second that were discarded before transmission because of a timeout condition. The timeout condition occurs while the switch or port waits for buffer credit from the receiving port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Credit Recovery Link Reset Rate	The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Discarded Frame Rate ¹	The average number of frames per second that are discarded because host buffers are unavailable for the port.
Link Reset Received Rate	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.
Link Reset Transmitted Rate	The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.
Port Congestion Index	The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Zero Buffer Credit Percentage	The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that metadata was collected. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.
Zero Buffer Credit Rate	The average number of Zero Buffer Credit conditions per second that occurred. A Zero Buffer Credit condition occurs when a port is unable to send frames because of a lack of buffer credit since the last node reset. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.
Notes:	
1. This metric is available only for ports on Brocade switches.	

Table 6. Link error rates

Metric	Description
Encoding Disparity	The average number of disparity errors per second that are received.
Invalid Link Transmission Rate	The average number of times per second that an invalid transmission word was detected by the port while the link did not experience any signal or synchronization loss.
Invalid Word Transmission Rate	The average number of bit errors per second that are detected.
Link Failure Rate	The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.
Loss of Signal Rate	The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.
Loss of Sync Rate	The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.
Primitive Sequence Protocol Error Rate	The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.

Table 7. Miscellaneous port metrics

Metric	Description
Link Quality Percentage	The percentage is based on whether the port is an expansion port (E_port) or a fabric port (F_port), and on the numbers and types of errors that are detected by the port.
Port Frame Size (Overall)	The average frame transfer size. This value is measured in KiB and includes frames that are sent and frames that are received by the port.
Port Frame Size (Receive)	The average size of a frame, in KiB, that is received by the port.
Port Frame Size (Send)	The average size of a frame, in KiB, that is sent through the port.

Monitoring the capacity of resources

IBM Spectrum® Control can collect information about the capacity and space usage of block, file, and object storage resources. This information includes key metrics that can help you measure, identify, and troubleshoot capacity and space usage issues in your storage. You can view capacity metrics for storage resources, including tiers.


You can use the capacity and space usage information to complete the following tasks:

- Measure, compare, and troubleshoot the capacity and space usage of resources such as:
 - Block storage systems and pools
 - Compressed, thin-provisioned, and Easy Tier® volumes

- File systems, pools, and filesets in file storage systems
- Containers in object storage systems
- Tiers
- Review the alerts that were triggered when the capacity or space usage of a resource fell outside a specific range.
- Customize views of capacity and space usage so that you can analyze specific resources and metrics during time ranges that you specify.
- View capacity and space usage information in a chart or table format to help you quickly identify where and when issues are occurring. The chart is a visual representation of how the capacity and space usage of resources trend over time.
- Drill down into resources to view detailed information about the capacity and space usage of internal resources. For example, if a SAN Volume Controller is shown in the chart, you can quickly view and compare the capacity and space usage of its pools.
- Implement server-centric monitoring of SAN resources without requiring a Storage Resource agent. When you add an agentless server, IBM Spectrum Control correlates the server with the known host connections on monitored storage systems. If matches are found, you can view the amount of storage that is assigned to the server, and trace that storage back to the storage system. You can then view details about the internal resources of the related storage system, including capacity and space usage information.
- Export capacity and space usage information to a CSV file. A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application.

Before you begin

Before you can view capacity and space usage information for resources, you must complete the following tasks:

Task	 Learn more
Add storage systems for monitoring by IBM Spectrum Control.	For information about how to add these resources, see Adding storage systems .
Ensure that capacity data is being collected for a resource. Typically, data collection is scheduled when resources are added for monitoring.	For information about verifying data collection, see Collecting data .
Define capacity and capacity usage alerts to be notified if the capacity or capacity usage of a resource falls outside a specified range and might represent a potential problem.	For information about how to define capacity and capacity usage alerts, see Defining alert definitions for general attributes and capacity changes .
When you define a capacity or capacity usage alert for an alert policy or for an individual resource, select a specific metric that you want to measure. For example, you can define an alert that notifies you when the total used capacity for a storage system falls outside a specified range.	

- [Setting capacity limits](#)
If your company has a policy to maintain a reserve of available capacity, you can set a limit on the amount of capacity that is used by your storage systems and pools.
- [Viewing capacity information](#)
View the capacity and space usage information of resources such as block storage systems, volumes, and pools. You can view similar information for filesets, file systems, and file system pools as internal resources of file storage systems. You can view information for containers as internal resources of object storage systems. Additionally, you can view capacity and space usage information for resources that violated a specified capacity range and generated an alert.
- [Viewing the capacity of external storage](#)
View the used capacity and capacity of external pool storage that is used by file systems in IBM Spectrum Scale. External pools can include storage that is provided by IBM® Cloud Object Storage, Amazon Simple Storage Service (S3), OpenStack Swift, IBM Spectrum Archive, IBM Spectrum Protect, and other storage providers.
- [Viewing capacity alerts and violations](#)
View the alerts that were triggered when the capacity of a resource changes and reaches a threshold. For example, you can view alerts that are generated when the measured value of a capacity metric for a pool meets the conditions for generating an alert.
- [Capacity metrics](#)
Use IBM Spectrum Control to collect and view capacity metrics about the storage systems in your environment.

Related reference

- [Capacity metrics for block storage systems](#)
- [Capacity metrics for file storage systems](#)
- [Capacity metrics for object storage systems](#)
- [Capacity metrics for tiers](#)

Setting capacity limits

If your company has a policy to maintain a reserve of available capacity, you can set a limit on the amount of capacity that is used by your storage systems and pools.

Before you begin

Only users with Administrator privileges can set capacity limits.

About this task

When you set a limit on the amount of capacity that is used, you can monitor whether the capacity that is used is above or below the limit that you set. You can also define alerts that notify you when you are over the capacity limit.

You can set the same capacity limit:

- For all storage systems or one or more storage systems
- For all pools or one or more pools

Alternatively, you can set different capacity limits for storage systems and pools.

Tip: FlashSystem A9000 and FlashSystem A9000R storage systems are configured to report storage consumption at the storage system rather than the pool level. So, if you want to set a capacity limit for these types of storage systems, set the capacity limit for the storage systems.

Procedure

1. To set the capacity limit for storage systems or pools, choose one of the following options:

Storage Resource	Step
Block Storage Systems	Click Storage, and then click Block Storage Systems.
Pools	Click Storage, and then click Pools.

2. Select the storage systems or pools.
Tip: To select multiple storage systems or pools, press Shift and click.
3. Right-click and click View Properties.
4. Click Edit Properties
5. Enter the percentage value for the capacity limit.
6. Save your changes.

Results

The capacity limit is set and can be shown as a percentage value or GiB value. Just right-click any column heading and add Capacity Limit (%) or Capacity Limit (GiB) to the Block Storage Systems or Pools pages. You can also add the Adjusted Used Capacity (%) column which tells you how much used capacity is left based on the capacity limit. By default, the Capacity-to-Limit (GiB) column is shown on the Block Storage Systems and Pools pages so that you can see how much available capacity you have left before you reach the capacity limit.

Zero capacity: When you set the capacity limit for pools, the values shown for Zero Capacity are readjusted to take into account the capacity limit of the pool. The date will represent when the capacity limit of the pool is reached. If the pool has already reached the capacity limit, Depleted is shown. None is shown when a trend in storage consumption can't be detected because the pool's storage isn't being consumed or because not enough data was collected to predict storage consumption.

- [Tutorial: Investigating compliance with the capacity limit](#)
Track the compliance of your block storage systems with the capacity limit.
- [Removing the capacity limit](#)
Remove the capacity limit that was set for your storage systems and pools.
- [Capacity limit metrics](#)
Monitor how much capacity is available for storing data when you set a capacity limit for your storage systems and pools.

Tutorial: Investigating compliance with the capacity limit

Track the compliance of your block storage systems with the capacity limit.

About this task


In this tutorial, you set the capacity limit for the storage systems on your production platform to 80%. Today, you want to create a chart that plots the compliance of your block storage systems with the capacity limit that was set.

Because you want to track capacity in relation to the capacity limit, you remove the Capacity (GiB) metric from the chart. And, you add the following metrics to track whether the available capacity of each storage system is over or under the capacity limit:

- Capacity Limit (GiB)
- Capacity-to-Limit (GiB)

Procedure

1. From the Storage menu, click Block Storage Systems.
2. Click the Capacity tab.
3. Select the storage systems that you want to add to the capacity chart.
To select multiple storage systems, press shift and click. You can add up to 10 storage systems at a time to the capacity chart.
4. Click the Select Chart Metrics button, add Capacity Limit (GiB) and Capacity-to-Limit (GiB), and click Save.



Tip: To see the capacity metrics for a single resource, click a storage system in the Selected Resources section of the chart. And, if you want to focus on a specific period, you can change the default date range of a month to a week, or select start and end dates.

Results

You can identify the storage systems that comply with the capacity limit that was set.

Removing the capacity limit

Remove the capacity limit that was set for your storage systems and pools.

About this task

Only users with Administrator privileges can remove the capacity limit.

Procedure

1. To remove the capacity limit for storage systems or pools, choose one of the following options:

Storage Resource	Step
Block Storage Systems	Click Storage, and then click Block Storage Systems.
Pools	Click Storage, and then click Pools.

2. Select the storage systems or pools.
Tip: To select multiple storage systems or pools, press shift and click.
3. Right-click, and then click View Properties.
4. Click Edit .
5. Complete one of the following action:
 - If the pools or storage systems have the same capacity limit, delete the capacity limit.
 - If the pools or storage systems have different capacity limits, enter 0 in the Capacity Limit (%) field and press the Backspace key.
6. Save your changes.

Capacity limit metrics

Monitor how much capacity is available for storing data when you set a capacity limit for your storage systems and pools.

Capacity Limit (%)

As an Administrator, you can set a Capacity Limit (%) for block storage systems, or pools, or both. In this scenario, you want to keep 20% of your usable capacity in reserve so you set the capacity limit to 80%.

Example: Administrator Sets Capacity Limit to 80%



To see the GiB or percentage value for the capacity limit, right-click any column heading on the Block Storage Systems or the Pools pages and add them.

To monitor how much capacity is available in relation to the capacity limit that you set, the used and available capacity values are adjusted and are shown in the following columns:


- Adjusted Used Capacity (%)
- Capacity-to-Limit (GiB)

To see the values for Adjusted Used Capacity (%), right-click any column heading on the Block Storage Systems page or Pools page and add it.

Adjusted Used Capacity (%)

To know how much capacity can be used without exceeding the capacity limit, monitor the value for Adjusted Used Capacity (%).

Example: Adjusted Used Capacity




The formula for calculating Adjusted Used Capacity (%) is (Used Capacity in GiB/Capacity Limit in GiB)*100. Using the values that are shown in the illustration above, the value is (40 GiB/80 GiB)*100 or 50%. So, in this scenario, you can use 30% or 40 GiB of the usable capacity of the resource before you reach the capacity limit.

If the used capacity exceeds the capacity limit, the value for Adjusted Used Capacity (%) is over 100%.

Capacity-to-Limit (GiB)

Use the value for Capacity-to-Limit (GiB) to monitor how much available capacity is left before the capacity limit is reached.

Example: Capacity-to-Limit



The formula for calculating Capacity-to-Limit (GiB) is (Capacity Limit in GiB - Used Capacity in GiB). Using the values that are shown in the illustration above, the value is (80 GiB - 40 GiB) or 40 GiB.

If the capacity limit is exceeded, a negative value is shown for Capacity-to-Limit (GiB).

Viewing capacity information

View the capacity and space usage information of resources such as block storage systems, volumes, and pools. You can view similar information for filesets, file systems, and file system pools as internal resources of file storage systems. You can view information for containers as internal resources of object storage systems. Additionally, you can view capacity and space usage information for resources that violated a specified capacity range and generated an alert.

Before you begin

Before you view the capacity and space usage of resources, ensure that capacity data was collected for those resources during the time ranges that you want to analyze. Data collection is usually scheduled when resources are added for monitoring by IBM Spectrum Control.

About this task

You can view the following capacity information:

- [Capacity charts for block storage](#)
- [Capacity charts for file storage](#)
- [Capacity charts for object storage](#)
- [Capacity charts for tiers](#)

Capacity charts for block storage

Table 1. View capacity charts for block storage

Actions	Navigation
View capacity charts for block storage systems.	<ol style="list-style-type: none">From the Storage menu, click Block Storage Systems.Right-click one or more storage systems and click View Capacity.
View the capacity charts for all block storage pools.	<ol style="list-style-type: none">From the Storage menu, click Pools.Right-click one or more pools and click View Capacity.
View the capacity chart for the pools in a storage system.	<ol style="list-style-type: none">From the Storage menu, click Block Storage Systems.Right-click the storage system and click View Details.In the navigation menu, click Pools.Right-click one or more pools and click View Capacity.
View the capacity chart for all of the compressed, thin-provisioned, and Easy Tier® volumes.	<ol style="list-style-type: none">From the Storage menu, click Volumes.Right-click one or more volumes that are compressed, or thin-provisioned, or enabled for Easy Tier, and then click View Capacity.
View the capacity chart for all of the compressed, thin-provisioned, and Easy Tier volumes in the storage system.	<ol style="list-style-type: none">From the Storage menu, click Block Storage Systems.Right-click the storage system and click View Details.In the navigation menu, click Volumes.Right-click one or more volumes that are compressed, or thin-provisioned, or enabled for Easy Tier, and then click View Capacity.

Capacity charts for file storage

Table 2. View capacity charts for file storage

Actions	Navigation
View capacity charts for all of the filesets in a file storage system.	<ol style="list-style-type: none">From the Storage menu, click File Storage Systems.Right-click the storage system and click View Details.In the navigation menu, click Filesets.Right-click one or more filesets and click View Capacity.

Actions	Navigation
View the capacity charts for all of the file systems in a file storage system.	<ol style="list-style-type: none"> 1. From the Storage menu, click File Storage Systems. 2. Right-click the storage system and click View Details. 3. In the navigation menu, click File Systems. 4. Right-click one or more file systems and click View Capacity.
View the capacity charts for all of the file system pools in a file storage system.	<ol style="list-style-type: none"> 1. From the Storage menu, click File Storage Systems. 2. Right-click the storage system and click View Details. 3. In the navigation menu, click File System Pools. 4. Right-click one or more file system pools and click View Capacity.

Capacity charts for object storage

Table 3. View capacity charts for object storage

Actions	Navigation
View the capacity charts for all of the containers in an object storage system.	<ol style="list-style-type: none"> 1. From the Storage menu, click Object Storage Systems. 2. Right-click the storage system and click View Details. 3. In the navigation menu, click Containers. 4. Right-click one or more containers and click View Capacity.

Capacity charts for tiers

Table 4. View capacity charts for tiers

Actions	Navigation
View the capacity charts for tiers.	<ol style="list-style-type: none"> 1. From the Groups menu, click Tiers. 2. Right-click one or more tiers and click View Capacity.

- [How capacity information is displayed](#)
The capacity view is displayed when you view the capacity of resources, such as block storage systems or pools, including the capacity of resources in alert violations. Information in the capacity view is organized into a chart or table and a legend.
- [Investigating capacity trends for block storage systems](#)
Use the capacity charts and the information that is shown in the tables for the block storage systems to check which storage systems have the highest growth rates and which storage systems might need more capacity.
- [Investigating capacity trends for block storage pools](#)
Use the capacity charts and the information that is shown in the tables for the block storage pools to check which storage pools have the highest growth rates, which storage pools might need more capacity, and which storage pools are at risk because too much capacity is committed to the thin-provisioned volumes.
- [Investigating capacity trends for volumes](#)
Use the capacity charts and the information that is shown in the tables to check the rate at which the thin-provisioned volumes consume capacity, the compression savings for the volumes that were converted to compressed volumes, and the changes in the distribution of Easy Tier volumes across the SSD, Enterprise HDD, and Nearline HDD drives.
- [Investigating capacity trends for file systems](#)
Use the capacity charts and the information that is shown in the tables for the file systems to check which file systems have the highest growth rates and which file systems might require additional capacity.
- [Investigating capacity trends for file system pools](#)
Use the capacity charts and the information that is shown in the tables for the file system pools in your storage systems to check which storage pools have the highest growth rates so that you can plan future storage requirements.
- [Investigating capacity trends for filesets](#)
Use the capacity charts to check which filesets have the highest growth rates in your file storage systems.
- [Investigating capacity trends for containers](#)
Use the capacity charts to check which containers have the highest growth rates in your object storage systems.
- [Investigating capacity trends for tiers](#)
Use the capacity charts and the information that is shown in the tables to check which tiers have the highest growth rates and which tiers might require more capacity.
- [Investigating capacity trends for servers](#)
Use the capacity charts and the information that is shown in the tables for the servers to check which servers have the highest growth rates for mapped and used SAN capacity. Use this information to help determine and compare the capacity usage rate of servers and which ones might need more capacity.
- [Creating bookmarks for URLs of capacity views](#)
You can create a bookmark for the URL of a capacity view. You can also open a duplicate of the capacity view to change and compare views.

Related reference

- [Capacity metrics for block storage systems](#)
- [Capacity metrics for file storage systems](#)
- [Capacity metrics for object storage systems](#)
- [Capacity metrics for tiers](#)

How capacity information is displayed

The capacity view is displayed when you view the capacity of resources, such as block storage systems or pools, including the capacity of resources in alert violations. Information in the capacity view is organized into a chart or table and a legend.

Capacity chart

The top section of the capacity view shows information about the selected resources, such as block storage systems or pools. You can view this information in the following formats:



Chart

The chart shows a visual representation of how the capacity of a resource trends over time. Each line on the chart represents a metric and a resource. For example, if you select two metrics and three resources, six lines are shown on the chart. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.



Table

The table shows capacity information that is formatted into rows and columns. Each row represents a resource and a time stamp from the chart; each column represents a metric from the chart. For example, if two metrics and three resources are displayed on a chart, and each of the lines on the chart have 10 data points, 30 rows and 2 columns are shown in the table.

To view other metrics and asset information for a resource, right-click anywhere in the header row of the table and select extra columns. The type of resource determines the metrics and information that is available.

Capacity chart legend

The bottom section of the capacity view shows more information about the resources in the chart. The information is formatted into rows and columns. Each row represents a resource that was selected for the view. Each column provides capacity and space usage information about a resource.

You can also view information about the resources that are related to the resources in the legend. For example, in the capacity view for block storage systems, if a SAN Volume Controller is listed in the legend, right-click it to view the capacity of its internal resources, such as pools.

Tip:

- If you hide a resource in the capacity chart, the row for that resource remains visible in the chart legend.
- In the chart legend, information about a resource represents the state of that resource during the selected time range. Specifically, this information does not represent the current state of a resource, but instead shows the final state of the resource during the selected time range. For example, if the current date is January 1, and you view a pool with the time range set to December 1 to December 7, the legend shows the capacity of that pool on December 7. If the capacity of the pool was changed between December 7 and January 1, this historical capacity is different from the current capacity of the pool on January 1.

- **[Controls for capacity views](#)**

Each capacity view includes controls for customizing how information is displayed. The type of view and the resources that you are viewing determine which controls are available.

- **[Resources in the capacity chart legend](#)**

In the capacity view, capacity metrics and related information for resources are shown in a chart and in the chart legend. For resources in the chart legend, such as block storage systems, you can open separate capacity views for their internal resources, such as pools or volumes. The separate capacity views use the same time range as the capacity view in the main window of the GUI.

Related reference

- [Capacity metrics](#)

Controls for capacity views

Each capacity view includes controls for customizing how information is displayed. The type of view and the resources that you are viewing determine which controls are available.



View chart

View capacity information in a chart format. The chart shows a visual representation of how resource capacity trends over time. Each line on the chart represents a metric and a resource. The y-axis shows the unit of measurement for a metric. If more metrics were selected with a different unit of measurement, an extra y-axis is shown on the right side of the chart window.

Hover the mouse pointer over points on a line to view a snapshot of capacity information at a specific time.



View table

View capacity information in a table format. Each row represents a resource and a time stamp. Each column represents a metric. You can view other metrics and information for a resource by right-clicking anywhere in the header row for the table and selecting more columns. The type of resource determines the metrics and information that are available.

For information how to filter, sort, and customize the columns in a table, see [Customizing lists](#).



Metrics

Add metrics to the capacity chart. On the Select Chart Metrics dialog, the metrics that are available depend on the type of resource that is being shown in the chart. The number next to the name of the metrics type represents the number of metrics that are currently selected.

You can select multiple metrics at the same time, but you cannot include more than two unit types in the same view. For example, if you select metrics that use % and ops/s as units of measurement, you cannot select more metrics that use different units of measurement such as KiB/op or MiB/s.

For a list of metrics that you can view for resources, see [Capacity metrics](#).



Hide and show resources

When the chart includes multiple resources, you can click the icon next to a resource to show only the line for that resource. Each icon is shown in a different color to match the color of line for the resource.

You can also show and hide resources in the chart by selecting resources in the chart legend. To select multiple resources at the same time, press Shift or Ctrl and click those resources. Press Shift and click to select consecutive rows in the legend; press Ctrl and click to select non-consecutive rows.

You can show up to 10 resources in a chart at the same time.

Specify a time range

The time range of the capacity information is shown below the chart. You can change this range to display information for different times when data was collected. When you first access the capacity view, the default time range is the last month.

Ensure that capacity data was collected during the time range that you select. If data was not collected during the time range, the chart and table are blank. If data collection was interrupted during the time range, the chart and table show gaps for the time increments when data was not collected.

For example, if you select a time range for the last 7 days, but data was not collected on days 4 and 5, the lines in the chart do not show data for days 4 and 5.



Export information about the chart to a file

Export information on a capacity view to a CSV file.



Open the capacity view in a separate web browser window

Open a duplicate of the current capacity view in a separate web browser window. You can change the information that is displayed in this separate window while retaining the original capacity view for comparison.

Actions for managing the resources in the chart legend

The chart legend in bottom section of the view shows more information about the selected resources. This information is organized into rows and columns, where each row represents a resource.

When you select one or more resources in the legend, the following actions are available in the Actions menu:

View Properties

View key details about a resource, including asset, status, configuration, and capacity information.

View *resource* Capacity

View the capacity of resources that are internal or related to a resource in the chart legend. For example, when you view the capacity of a SAN Volume Controller, you can right-click it and view the capacity of its internal resources, such as pools.

Information about an internal or related resource is shown in a separate web browser window. This window uses the same time range as the capacity view in the main window of the GUI.

Resources in the capacity chart legend

In the capacity view, capacity metrics and related information for resources are shown in a chart and in the chart legend. For resources in the chart legend, such as block storage systems, you can open separate capacity views for their internal resources, such as pools or volumes. The separate capacity views use the same time range as the capacity view in the main window of the GUI.

For example, when you view the capacity of a SAN Volume Controller, you can right-click it in the chart legend to open a capacity view for its internal resources, such as pools or volumes.

You can view capacity metrics for the following resources:

- Block storage systems
- Internal resources of block storage systems: volumes and pools
- Internal resources of file storage systems: filesets, file systems, and file system pools
- Internal resources of object storage systems: containers
- Tiers

The following table summarizes the resources and the internal resources for which you can view capacity metrics.

Table 1. Internal resources for which you can view capacity metrics

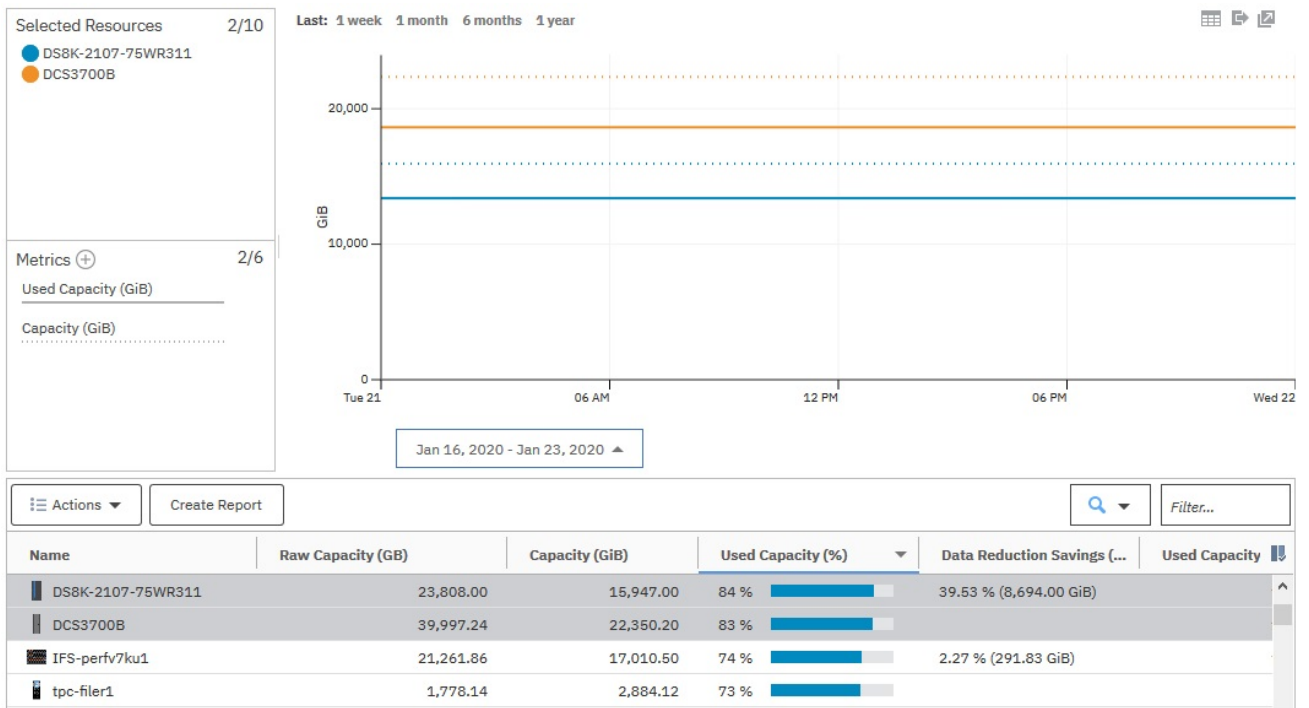
Resource	Internal resources
Block storage system	<ul style="list-style-type: none">• Volumes• Pools
File storage system	<ul style="list-style-type: none">• Filesets• File systems• File system pools
Object storage system	<ul style="list-style-type: none">• Containers

Investigating capacity trends for block storage systems

Use the capacity charts and the information that is shown in the tables for the block storage systems to check which storage systems have the highest growth rates and which storage systems might need more capacity.

Procedure

1. From the Storage menu, click Block Storage Systems.
2. Click Capacity and select the storage systems and the capacity metrics that you want to review.



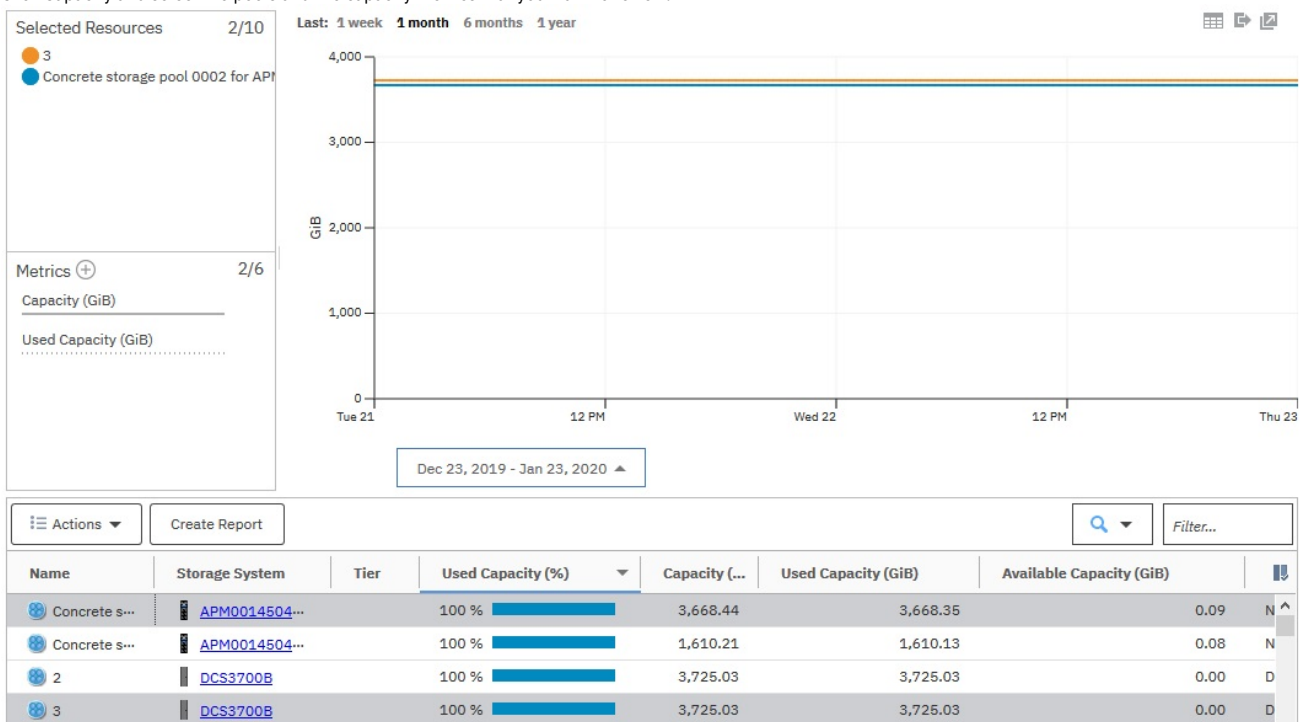
Tip: To check which pools consume the most storage, right-click the storage systems and click Capacity.

Investigating capacity trends for block storage pools

Use the capacity charts and the information that is shown in the tables for the block storage pools to check which storage pools have the highest growth rates, which storage pools might need more capacity, and which storage pools are at risk because too much capacity is committed to the thin-provisioned volumes.

Procedure

1. From the Storage menu, click Block Storage Systems.
Alternatively, you can click Storage > Pools.
2. Right-click the storage system and click View Details.
3. In the navigation pane, click Pools.
4. Click Capacity and select the pools and the capacity metrics that you want to review.




Tip: To investigate the pools with the highest shortfall rates, right-click the pools in the table and click Volume Capacity.

Investigating capacity trends for volumes

Use the capacity charts and the information that is shown in the tables to check the rate at which the thin-provisioned volumes consume capacity, the compression savings for the volumes that were converted to compressed volumes, and the changes in the distribution of Easy Tier® volumes across the SSD, Enterprise HDD, and Nearline HDD drives.

Procedure

- 1. From the Storage menu, click Block Storage Systems.
- 2. Right-click the storage system and click View Details.
- 3. In the navigation pane, click Volumes.
- 4. Click Capacity and select the volumes and the capacity metrics that you want to review.




Tip: To check the distribution of Easy Tier volumes, click Select Chart Metrics and add the Easy Tier metrics to the chart.

Investigating capacity trends for file systems

Use the capacity charts and the information that is shown in the tables for the file systems to check which file systems have the highest growth rates and which file systems might require additional capacity.

Procedure

- 1. From the Storage menu, click File Storage Systems.
- 2. Right-click the storage system and click View Details.
- 3. In the navigation menu, click File Systems.
- 4. Click Capacity and select the file systems and the capacity metrics that you want to review.



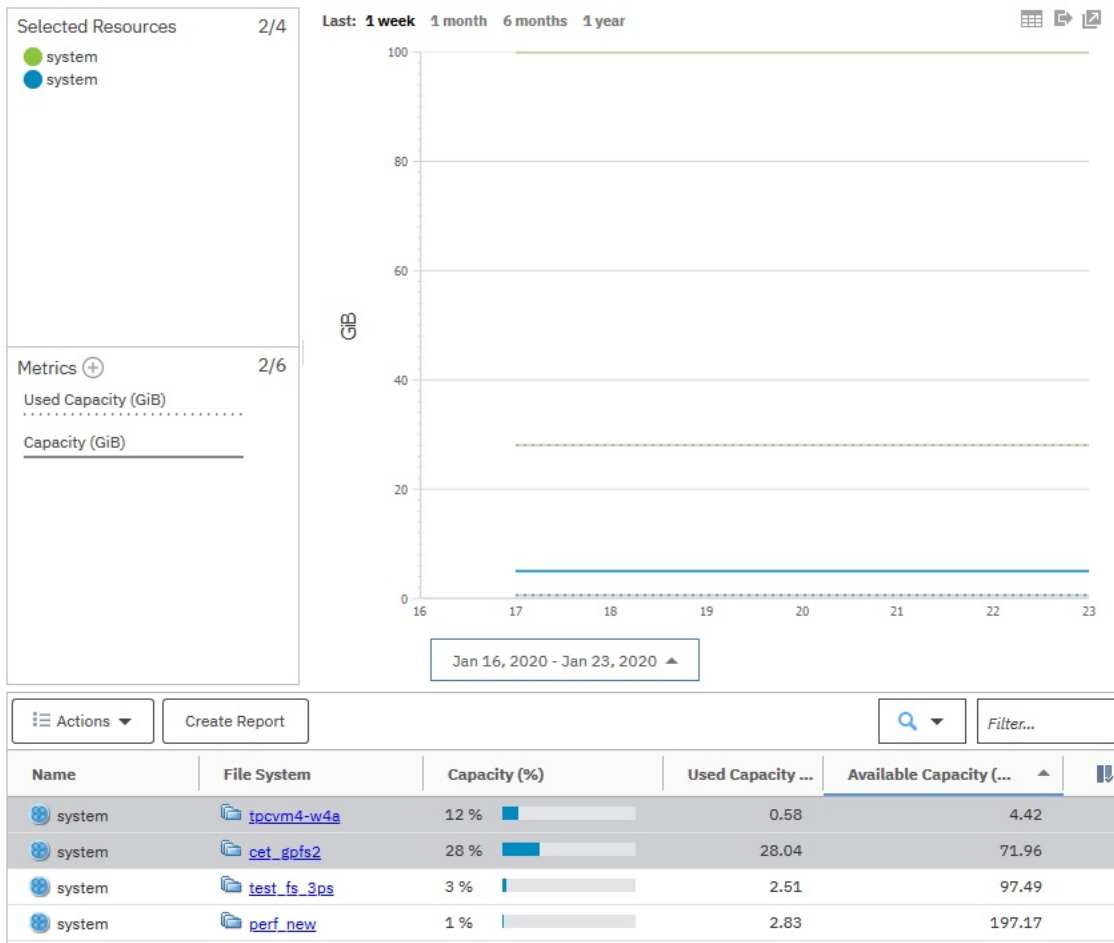
Tip: To check which pools consume the most storage, right-click the file systems and click File System Pool Capacity.

Investigating capacity trends for file system pools

Use the capacity charts and the information that is shown in the tables for the file system pools in your storage systems to check which storage pools have the highest growth rates so that you can plan future storage requirements.

Procedure

1. From the Storage menu, click File Storage Systems.
2. Right-click the storage system and click View Details.
3. In the navigation menu, click File System Pools.
4. Click Capacity and select the pools and the capacity metrics that you want to review.

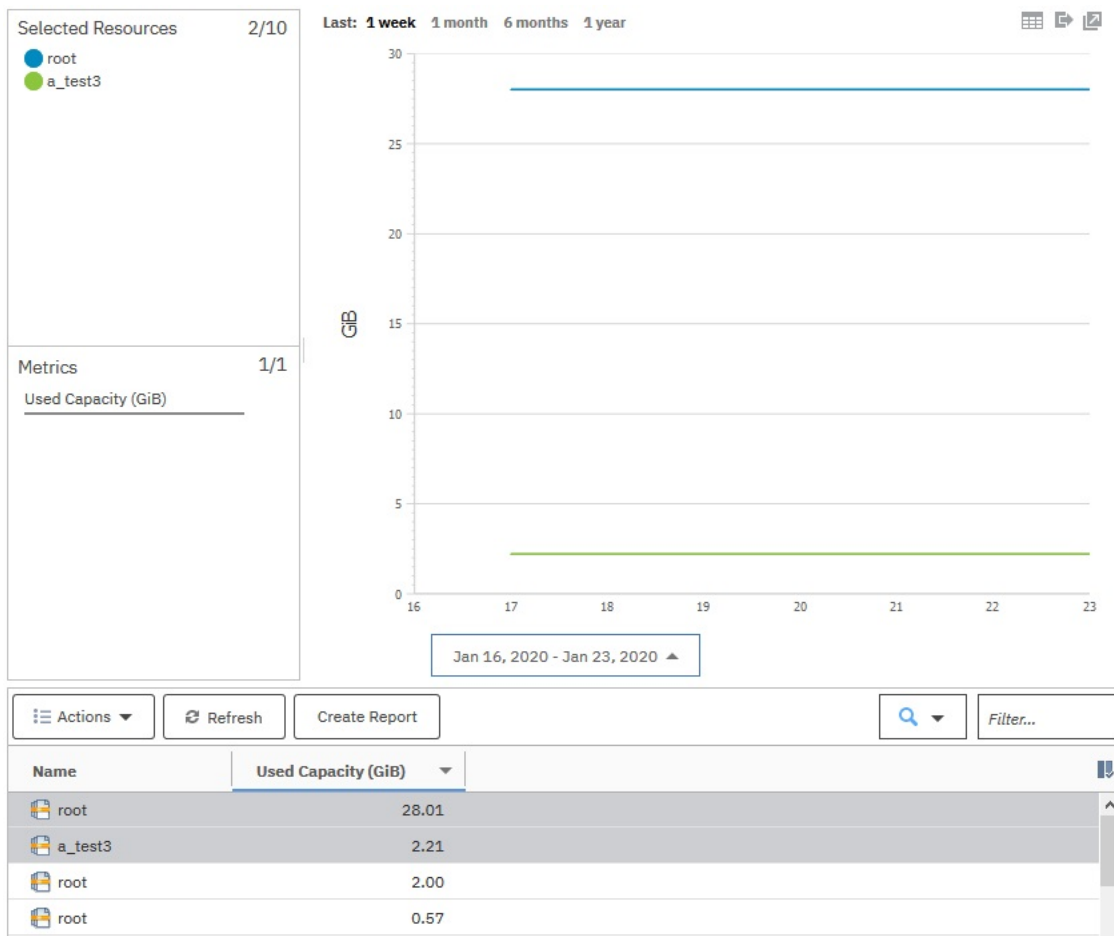


Investigating capacity trends for filesets

Use the capacity charts to check which filesets have the highest growth rates in your file storage systems.

Procedure

1. From the Storage menu, click File Storage Systems.
2. Right-click the storage system and click View Details.
3. In the navigation menu, click Filesets.
4. Click Capacity and select the filesets that you want to review.




Investigating capacity trends for containers

Use the capacity charts to check which containers have the highest growth rates in your object storage systems.

Procedure

1. From the Storage menu, click Object Storage Systems.
2. Right-click the storage system and click View Details.
3. In the navigation menu, click Containers.
4. Click Capacity and select the objects that you want to review.



Investigating capacity trends for tiers

Use the capacity charts and the information that is shown in the tables to check which tiers have the highest growth rates and which tiers might require more capacity.

Procedure

1. From the Groups menu, click Tiers.
2. Click Capacity.

Investigating capacity trends for servers

Use the capacity charts and the information that is shown in the tables for the servers to check which servers have the highest growth rates for mapped and used SAN capacity. Use this information to help determine and compare the capacity usage rate of servers and which ones might need more capacity.

Procedure

1. From the Servers menu, click Servers.
2. Click Capacity and select the servers and the capacity metrics that you want to review.

Creating bookmarks for URLs of capacity views

You can create a bookmark for the URL of a capacity view. You can also open a duplicate of the capacity view to change and compare views.

About this task

The bookmark stores details of the capacity view, such as the resources, metrics, interval, and other details. It also stores the time period that is specified in the time selector at the top of the chart.

Procedure

To create a bookmark for a capacity view, follow these steps:

1. Open the capacity view for which you want to create a bookmark.
2. Right-click and click the relevant option for creating a bookmark.
Tip: To open a duplicate of the current capacity view, click the icon for opening a new browser window. You can change the information that is displayed in the new window while retaining the original capacity view for comparison.

Viewing the capacity of external storage

View the used capacity and capacity of external pool storage that is used by file systems in IBM Spectrum Scale. External pools can include storage that is provided by IBM® Cloud Object Storage, Amazon Simple Storage Service (S3), OpenStack Swift, IBM Spectrum Archive, IBM Spectrum Protect, and other storage providers.

About this task


For all external storage that is being used by IBM Spectrum Scale file systems, you can complete the following tasks:

- View used capacity to understand how much data is being migrated from those file systems to external storage.
- Charge back storage costs to the departments within your organization.
- Gauge shortfall in case you need to recall data from external pools.

For external storage that is provided by IBM Cloud Object Storage, you can also view capacity information, including the percentage of space (active and inactive) that is being used. Use this information to complete the following tasks:

- View how storage space is distributed between internal and external storage tiers.
- Identify the external pools that are running out of space.
- Determine when you need to reconcile files between an IBM Spectrum Scale and an external storage tier. For example, when you remove files from a file system, no policy is available to automatically remove the cloud objects that are created. This situation might cause orphan objects (inactive data) on the cloud. To clean up inactive data and keep the cloud in sync with the file system, you can reconcile files.

The following figure illustrates the relationship between IBM Spectrum Scale and the IBM Cloud Object Storage system that provides storage for external pools:



Procedure

To view the used capacity and capacity of external storage, complete the following steps:

1. From the Storage menu, click File Storage Systems.
2. Right-click a IBM Spectrum Scale storage system and click View Details.
3. In the navigation pane, click Pools.
4. View the following columns for information about the used capacity and capacity of external pools:

Capacity (%)

The percentage of used capacity in the file system pool. Hover over the bar for a specific pool to view statistics about its used and available capacity.

For external pools that are connected to cloud services or storage providers, only the amount of used capacity is shown. Cloud services and other storage providers might include IBM Spectrum Protect, IBM Spectrum Archive, Amazon Simple Storage Service (S3), or OpenStack Swift.

For external pools that are connected to vaults in IBM Cloud Object Storage, the percentage of used capacity includes active and inactive data. Space values are determined from the usable vault capacity that is associated with the cloud account and the hard quota value (if a hard quota is configured). Hover over the bar for a specific pool to view statistics about its used active space, used inactive space, and available capacity.

Active and inactive data in external pools

Active data is data that has corresponding stub files on the GPFS file system. Inactive data is data that exists in the external pool but is not accessible in the GPFS file system. Data is classified as inactive under these conditions:

- When you migrate a file off a GPFS file system, its data is still visible and accessible through its stub file. When the stub file is deleted, the data in the external pool is no longer accessible and is considered inactive.
- When you modify a migrated file, it's recalled from the external pool and modified locally. The previous data for the file remains in the external pool but is no longer accessible through the GPFS file system and is considered inactive.

Tip: To keep the files in sync between an external pool and a GPFS file system and clean up inactive data, you can run the IBM Spectrum Scale `mmcloudgateway files reconcile` command. For more information, see [Reconciling files between IBM Spectrum Scale file system and cloud storage tier](#).

External Used Capacity (GiB)

The amount of used capacity on an external pool. This value includes active data and inactive data. Active data is data that has corresponding stub files on the GPFS file system. Inactive data is data that exists in the external pool but is not accessible in the GPFS file system.

Availability: External pools that are provided by IBM Cloud Object Storage.

Inactive Used Capacity

The amount of inactive used capacity on an external pool. Inactive data is data that exists in the external pool but is not accessible in the GPFS file system.

Available for external pools that are provided by Cloud Object Storage only.

Availability: External pools that are provided by IBM Cloud Object Storage.

Used Capacity

The amount of capacity in the file system pool that is being used. For external pools, used capacity includes active data only. Active data is data that has corresponding stub files on the GPFS file system.

Availability: All storage systems.

- **Identifying shortfall before data is recalled from external storage**

View the percentage of migrated data in external pools that does not fit into the available capacity on an IBM Spectrum Scale file system.

Related information

- [Introduction to Transparent Cloud Tiering for IBM Spectrum Scale](#)

Identifying shortfall before data is recalled from external storage

View the percentage of migrated data in external pools that does not fit into the available capacity on an IBM Spectrum Scale file system.

About this task

This percentage, called *shortfall*, represents the relative risk of not having enough internal file system space if you recall all data from the associated external pools. The higher the percentage, the higher the amount of migrated data that does not fit.

Procedure

To view the potential shortfall for a file system, complete the following steps:

1. From the Storage menu, click File Storage Systems.
2. Right-click a IBM Spectrum Scale storage system and click View Details.
3. Under Internal Resources, click File Systems.
4. Locate the file system that you want to check and view the value in the Shortfall column.

To calculate shortfall, the following formula is used:

$$[(\text{External Pool Migrated Data} - \text{Available Capacity}) \div \text{External Pool Migrated Data}] \times 100$$

For example, if the available capacity on a file system is 200 GiB, but the used capacity of migrated data on external pools is 500 GiB, the shortfall percentage is 60% (300 GiB). In this case, consider adding capacity to the file system before recalling data.

$$[(500 - 200) \div 500] \times 100 = 60\%$$

If the file system has sufficient available capacity to contain all the migrated data in external pools, no shortfall exists. For example, if the available capacity on a file system is 500 GiB, and the used capacity in external pools is only 200 GiB, the shortfall is 0%.

Tip: Only the space that is used by migrated data in external pools is included in the calculation for shortfall. Pre-migrated data is not included because it exists in both the internal and external pools.

Viewing capacity alerts and violations

View the alerts that were triggered when the capacity of a resource changes and reaches a threshold. For example, you can view alerts that are generated when the measured value of a capacity metric for a pool meets the conditions for generating an alert.

Before you begin

To view alert violations for the capacity of a resource, you must collect capacity data and define capacity alerts for that resource or for the alert policy that manages the resource.

Procedure

1. For block storage systems, go to Storage > Block Storage Systems. For file storage systems, go to Storage > File Storage Systems.
2. Right-click a storage system and click View Details.
3. In the navigation menu, click Alerts.
4. Right-click the resource alert that you want to view, such as an alert that was defined for a capacity metric for the resource, and click View Alert.

Related concepts

- [Collecting data](#)
- [Alerting](#)

Related tasks

- [Defining alert definitions for general attributes and capacity changes](#)

Capacity metrics

Use IBM Spectrum® Control to collect and view capacity metrics about the storage systems in your environment.

- [Capacity metrics for block storage systems](#)
To review trends in capacity and space usage for storage, you add metrics to capacity charts. You use the charts to detect capacity shortages and space usage trends.

- [Capacity metrics for file storage systems](#)
To review trends in capacity and space usage for file storage systems, you add metrics to capacity charts. You use the charts for filesets, file systems, and file system pools to detect capacity shortages and space usage trends.
- [Capacity metrics for object storage systems](#)
To review trends in capacity and space usage for object storage systems, you add metrics to capacity charts. Use the charts for containers to detect capacity shortages and space usage trends for the containers in your object storage systems.
- [Capacity metrics for tiers](#)
Review trends in the capacity and space usage for tiers. You use the charts to detect capacity shortages and space usage for the tiers in your storage environment.

Capacity metrics for block storage systems

To review trends in capacity and space usage for storage, you add metrics to capacity charts. You use the charts to detect capacity shortages and space usage trends.

Alphabetical lists of the capacity and space usage metrics that you can add to charts are provided in the following sections:

- [Storage system capacity metrics](#)
- [Pool capacity metrics](#)
- [Volume capacity metrics](#)

Tip: IBM Spectrum® Control displays capacity values in base 2 (GiB), while the XIV® management GUI displays capacity values in base 10 (GB, TB). Even though different units of measurement are used, the storage values are equivalent. For more information about units of measurement, see [Units of measurement for storage data](#).

Storage system capacity metrics

To detect capacity shortages and investigate space usage trends, you can add the following metrics to the capacity chart for storage systems:

Adjusted Used Capacity (%)

The amount of capacity that can be used without exceeding the capacity limit.


Example: Adjusted Used Capacity

Before Capacity Limit Was Set



■ Used Capacity = 40 GiB
■ Available Capacity = 60 GiB

After Capacity Limit Was Set



— Capacity Limit = 80% or 80 GiB
■ Adjusted Used Capacity = 50% or 40 GiB
└─ Capacity-to-Limit = 30% or 40 GiB

The formula for calculating Adjusted Used Capacity (%) is $(\text{Used Capacity in GiB} / \text{Capacity Limit in GiB}) * 100$. For example, if the capacity is 100 GiB, the used capacity is 40 GiB, and the capacity limit is 80% or 80 GiB, then the value for Adjusted Used Capacity (%) is $(40 \text{ GiB} / 80 \text{ GiB}) * 100$ or 50%. So, in this example, you can use 30% or 40 GiB of the usable capacity of the resource before you reach the capacity limit.

If the used capacity exceeds the capacity limit, the value for Adjusted Used Capacity (%) is over 100%.

To add the Adjusted Used Capacity (%) column, right-click any column heading on the Block Storage Systems page.

See these related values for more information Capacity Limit (%), and Capacity-to-Limit (GiB).

This metric is not available for all storage systems, such as Dell EMC VMAX.

Available Capacity (GiB)

(Previously known as Available Pool Space) The total amount of the space in the pools that is not used by the volumes in the pools. To calculate available capacity, the following formula is used:

(pool capacity - used capacity)

For XIV systems, pool capacity is the physical capacity of the pools and does not include the provisioned capacity of the pools.

Availability: All storage systems.

Available Volume Capacity (GiB)

(Previously known as Effective Unallocated Volume Space) The total amount of remaining space that can be used by the volumes in the pools. The following formula is used to calculate this value:

[Provisioned Capacity - Used Capacity]

The capacity that is used by thin-provisioned volumes is typically less than their provisioned capacity. Therefore, the available capacity represents the difference between the provisioned capacity and the used capacity for all the volumes in the pools. For Hitachi VSP non-thin provisioned pool capacity, the available capacity is always zero.

Availability: All storage systems.

Capacity (GiB)

(Previously known as Pool Capacity) The total amount of storage space in the pools. For XIV systems and IBM Spectrum Accelerate, capacity represents the physical ("hard") capacity of the pool, not the provisioned ("soft") capacity. Pools that are allocated from other pools are not included in the total pool space.

Availability: All storage systems.

Capacity Limit (%) and Capacity Limit (GiB)

The limit that was set on the capacity that is used by your storage systems. For example, the policy of your company is to keep 20% of the usable capacity of your storage systems in reserve. So, you log into the GUI as Administrator and set the capacity limit to 80%.

Example: Administrator Sets Capacity Limit to 80%



Click the illustration above to find out how to set capacity limits.

The GiB value for the capacity limit for the storage system is calculated when you set the value for the Capacity Limit (%).

To add the Capacity Limit (%) and the Capacity Limit (GiB) columns, right-click any column heading on the Block Storage Systems page.


See these related values for more information Adjusted Used Capacity (%) and Capacity-to-Limit (GiB).

This metric is not available for all storage systems, such as Dell EMC VMAX.

Capacity-to-Limit (GiB)

The amount of capacity that is available before the capacity limit is reached.

Example: Capacity-to-Limit



The formula for calculating Capacity-to-Limit (GiB) is (Capacity Limit in GiB - Used Capacity in GiB). For example, if the capacity limit is 80% or 80 GiB and the used capacity is 40 GiB, then the value for Capacity-to-Limit (GiB) is (80 GiB - 40 GiB or 80% - 50%) which is 30% or 40 GiB.

See these related values for more information Capacity Limit (%) and Adjusted Used Capacity (%).

This metric is not available for all storage systems, such as FlashSystem A9000, FlashSystem A9000R, and Dell EMC VMAX.

Compression Savings (%)

The estimated amount and percentage of capacity that is saved by using data compression, across all pools on the storage system. The percentage is calculated across all compressed volumes in the pools and does not include the capacity of non-compressed volumes.

For storage systems with drives that use inline data compression technology, the Compression Savings does not include the capacity savings that are achieved at the drive level.

The following formula is used to calculate the amount of storage space that is saved:

$$\text{written capacity} - \text{compressed size}$$

The following formula is used to calculate the percentage of capacity that is saved:

$$((\text{written capacity} - \text{compressed size}) \div \text{written capacity}) \times 100$$

For example, the written capacity, which is the amount of data that is written to the volumes before compression, is 40 GiB. The compressed size, which reflects the size of compressed data that is written to disk, is just 10 GiB. Therefore, the compression savings percentage across all compressed volumes is 75%.

Availability: FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.

For FlashSystem A9000 and FlashSystem A9000R, all volumes in the pools are compressed.

Deduplication Savings (%)

The estimated amount and percentage of capacity that is saved by using data deduplication, across all data reduction pools on the storage system. The percentage is calculated across all deduplicated volumes in the pools and does not include the capacity of volumes that are not deduplicated.

The following formula is used to calculate the amount of storage space that is saved:

$$\text{written capacity} - \text{deduplicated size}$$

The following formula is used to calculate the percentage of capacity that is saved:

$$((\text{written capacity} - \text{deduplicated size}) \div \text{written capacity}) \times 100$$

For example, the written capacity, which is the amount of data that is written to the volumes before deduplication, is 40 GiB. The deduplicated size, which reflects the size of deduplicated data that is written to disk, is just 10 GB. Therefore, data deduplication reduced the size of the data that is written by 75%.

This metric is available for FlashSystem A9000, FlashSystem A9000R, and resources that run IBM Spectrum Virtualize 8.1.3 or later.

Drive Compression Savings (%)

The amount and percentage of capacity that is saved with drives that use inline data compression technology. The percentage is calculated across all compressed drives in the pools.

The amount of storage space that is saved is the sum of drive compression savings.

The following formula is used to calculate the percentage of capacity that is saved:

$$((\text{used written capacity} - \text{compressed size}) \div \text{used written capacity}) \times 100$$

Availability: Storage systems that contain IBM FlashCore® Modules with hardware compression.

Mapped Capacity (GiB)

(Previously known as Assigned Volume Space) The total volume space in the storage system that is mapped or assigned to host systems, including child pool capacity.

Availability: All storage systems.

Overprovisioned Capacity (GiB)

(Previously known as Unallocatable Volume Space) The capacity that cannot be used by volumes because the physical capacity of the pools cannot meet the demands for provisioned capacity. The following formula is used to calculate this value:

$$[\text{Provisioned Capacity} - \text{Capacity}]$$

Availability: All storage systems.

Shortfall (%)

The percentage of space that is over committed to the pools with thin-provisioned volumes. For example, you commit 100 GiB of space to a thin-provisioned volume in a pool with a capacity of 50 GiB. As the space is used by the thin-provisioned volume in increments of 10 GiB, the space available for allocation decreases and the shortfall in capacity becomes more acute.

To calculate the shortfall, the following formula is used:

$$[(\text{overprovisioned capacity} \div \text{committed but available capacity}) \times 100]$$

A shortfall occurs when you commit more space to the volumes in the pools than is physically available to the pools. If the physical space available to the pools is less than the committed provisioned capacity, then the pools do not have enough space to fulfill the commitment to the provisioned capacity.

For example, the physical capacity of the pools is 70 GiB, but 150 GiB of provisioned capacity was committed to the thin-provisioned volumes. If the volumes are using 50 GiB, 100 GiB is committed to those volumes (150 GiB – 50 GiB) with only 20 GiB of available pool space (70 GiB – 50 GiB). Because only 20 GiB of the pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% [(80 GiB ÷ 100 GiB × 100)].

Availability: DS8000®, FlashSystem storage systems, Dell EMC VMAX, VNX, and VNXe storage systems, Hitachi Virtual Storage Platform, SAN Volume Controller, Storwize® family storage systems that are configured with block storage, XIV systems, and IBM Spectrum Accelerate storage systems.

Provisioned Capacity (%)

(Previously known as Virtual Allocation) The percentage of the physical capacity that is committed to the provisioned capacity of the volumes in the pools. If the value exceeds 100%, the physical capacity doesn't meet the demands for provisioned capacity.

To calculate provisioned capacity percentage, the following formula is used:

$$[(\text{provisioned capacity} \div \text{pool capacity}) \times 100]$$

For example, if the provisioned capacity percentage is 200% for a storage pool with a physical capacity of 15 GiB, then the provisioned capacity that is committed to the volumes in the pools is 30 GiB. Twice as much space is committed to the pools than is physically available to the pools. If the provisioned capacity percentage is 100% and the physical capacity is 15 GiB, then the provisioned capacity that is committed to the pools is 15 GiB. The total physical capacity that is available to the pools is used by the volumes in the pools.

A provisioned capacity percentage that is higher than 100% is considered to be aggressive because insufficient physical capacity is available to the pools to satisfy the allocation of the committed space to the compressed and thin-provisioned volumes in the pools. In such cases, you can check the Shortfall (%) value to determine how critical the shortage of space is for the storage system pools.

Availability: All storage systems.

Provisioned Capacity (GiB)

(Previously known as Total Volume Capacity)

The total amount of provisioned capacity of volumes within the pool. If the pool is a parent pool, it also includes the storage space that can be made available to the volumes in the child pools.

Availability: All storage systems.

Safeguarded Capacity (GiB)

The total amount of capacity that is used to store volume backups that are created by the Safeguarded Copy feature in DS8000.

Total Capacity Savings (%)

The estimated amount and percentage of capacity that is saved by using data deduplication, pool compression, thin provisioning, and drive compression, across all volumes in the pool.

The following formula is used to calculate the amount of storage space that is saved:

$$[\text{Provisioned Capacity} - \text{Used Capacity}]$$

The following formula is used to calculate the percentage of capacity that is saved:

$$((\text{Provisioned Capacity} - \text{Used Capacity}) \div \text{Provisioned Capacity}) \times 100$$

This metric is available for FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and storage systems that run IBM Spectrum Virtualize.

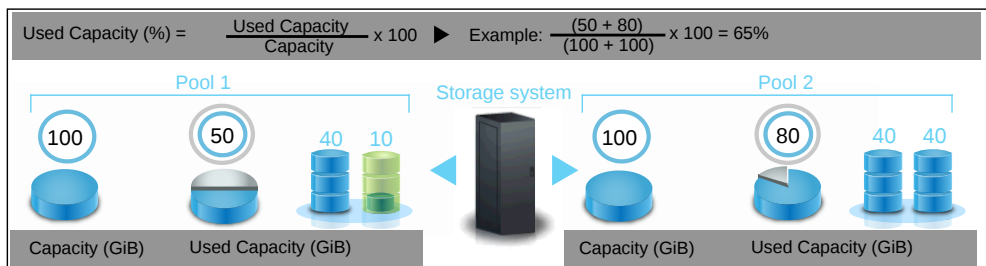
Unmapped Capacity (GiB)

(Previously known as Unassigned Volume Space) The total amount of space in the volumes that are not assigned to hosts.

Availability: All storage systems.

Used Capacity (%)

(Previously known as Physical Allocation)



The percentage of physical capacity in the pools that is used by the standard-provisioned volumes, the thin-provisioned volumes, and the volumes in child pools. Check the value for used capacity percentage to see:

- Whether the physical capacity of the pools is fully allocated. That is, the value for used capacity is 100%.
- Whether you have sufficient capacity to provision new volumes with storage
- Whether you have sufficient capacity to allocate to the compressed and thin-provisioned volumes in the pools

Availability: All storage systems.

Used Capacity (GiB)

(Previously known as Allocated Space) The amount of space that is used by the standard- and thin-provisioned volumes in the pools. If the pool is a parent pool, the amount of space that is used by the volumes in the child pools is also calculated.

The capacity that is used by for thin-provisioned volumes is less than their provisioned capacity, which is shown in the Provisioned Capacity (GiB) column. If a pool doesn't have thin-provisioned volumes, the value for used capacity is the same as the value for provisioned capacity.

Availability: All storage systems.

Pool capacity metrics

If sufficient data is collected, you can view charts that compare the capacity, used capacity, and available capacity of the pools in your data center.

In the Zero Capacity column on the Pools page, you can see the date, based on the storage usage trends for the pool, when the pool will run out of available capacity.

Zero Capacity: The capacity information that is collected over 180 days is analyzed to determine, based on historical storage consumption, when the pools will run out of capacity. The pools that have already run out of capacity are marked as depleted. For the other pools, a date is provided so that you know when the pools are projected to run out of capacity. If sufficient information isn't collected to analyze the storage usage of the pool, *None* is shown as the value for zero capacity. If a capacity limit is set for the pool, the date shown in the Zero Capacity column is the date when the available capacity based on the capacity limit will be depleted. For example, if the capacity limit for a 100 GiB pool is 80%, it is the date when the available capacity of the pool is less than 20 GiB. *Depleted* is shown in the column when the capacity limit is reached.

To detect capacity shortages and investigate trends in storage usage, you can add the following metrics to the capacity chart for pools:


Adjusted Used Capacity (%)

The amount of capacity that can be used without exceeding the capacity limit.

Example: Adjusted Used Capacity

Before Capacity Limit Was Set

100 GiB Capacity



■ Used Capacity = 40 GiB

■ Available Capacity = 60 GiB

After Capacity Limit Was Set

100 GiB Capacity



— Capacity Limit = 80% or 80 GiB

■ Adjusted Used Capacity = 50% or 40 GiB

▨ Capacity-to-Limit = 30% or 40 GiB

The formula for calculating Adjusted Used Capacity (%) is (Used Capacity in GiB/Capacity Limit in GiB)*100. For example, if the capacity is 100 GiB, the used capacity is 40 GiB, and the capacity limit is 80% or 80 GiB, then the value for Adjusted Used Capacity (%) is (40 GiB/80 GiB)*100 or 50%. So, in this example, you can use 30% or 40 GiB of the usable capacity of the resource before you reach the capacity limit.

If the used capacity exceeds the capacity limit, the value for Adjusted Used Capacity (%) is over 100%.

To add the Adjusted Used Capacity (%) column, right-click any column heading on the Pools page.

See these related values for more information Capacity Limit (%) and Capacity-to-Limit (GiB).

Availability: This metric is not available for all storage systems, such as Dell EMC VMAX.

Available Capacity (GiB)

(Previously known as Available Pool Space) The amount of physical space that is available in the pool. If the pool is a parent pool, the amount of space that is used by the volumes in the child pools is also included.

Availability: All storage systems. For FlashSystem A9000 and FlashSystem A9000R, this value represents provisioned capacity rather than physical space.

Available Repository Capacity (GiB)

The available, unallocated storage space in the repository for Track Space-Efficient (TSE) thin-provisioning.

Availability: DS8000 thin-provisioned pools.

Available Soft Capacity (GiB)

The amount of virtual storage space that is available to allocate to volumes in a storage pool.

Availability: XIV systems, and IBM Spectrum Accelerate storage systems.

Available Volume Capacity (GiB)

(Previously known as Effective Unallocated Volume Space) The total amount of remaining capacity that can be used by the existing volumes in the pools. The following formula is used to calculate this value:

Provisioned Capacity - Used Capacity

The capacity that is used by thin-provisioned volumes is typically less than their provisioned capacity. Therefore, the available capacity represents the difference between the provisioned capacity and the used capacity for all the volumes in the pool. For Hitachi VSP non-thin provisioned pool capacity, the unused volume capacity is always zero.

Availability: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Available Written Capacity (GiB)

(Previously known as Effective Used Capacity) The amount of capacity that can be written to the pools before inline compression is applied. If the pools are not compressed, this value is the same as Available Capacity.

Important: Because data compression is very efficient, a pool can run out of Available Written Capacity while physical capacity is still available. To stay aware of your capacity needs, monitor this value and Available Capacity.

Capacity (GiB)

The total amount of storage space in the pool. For XIV systems and IBM Spectrum Accelerate, capacity represents the physical or ("hard") capacity of the pool, not the provisioned ("soft") capacity.

Availability: All storage systems.

Capacity Limit (%) and Capacity Limit (GiB)

The limit that was set on the capacity that is used by your pools. For example, the policy of your company is to keep 20% of the usable capacity of your pools in reserve. So, you log into the GUI as Administrator and set the capacity limit of your pools to 80%.

Example: Administrator Sets Capacity Limit to 80%



Click the illustration above to find out how to set capacity limits.

The GiB value for the capacity limit for the pool is calculated when you set the value for the Capacity Limit (%).

To add the Capacity Limit (%) and the Capacity Limit (GiB) columns, right-click any column heading on the Pools page.

See these related values for more information Adjusted Used Capacity (%) and Capacity-to-Limit (GiB).


Zero capacity: When you set the capacity limit for pools, the values shown for Zero Capacity are readjusted to take into account the capacity limit of the pool. The date will represent when the capacity limit of the pool is reached. If the pool has already reached the capacity limit, Depleted is shown. None is shown when a trend in storage consumption can't be detected because the pool's storage isn't being consumed or because not enough data was collected to predict storage consumption.

Availability: This metric is not available for all storage systems, such as Dell EMC VMAX.

Capacity-to-Limit (GiB)

The amount of capacity that is available before the capacity limit is reached.

Example: Capacity-to-Limit



The formula for calculating Capacity-to-Limit (GiB) is (Capacity Limit in GiB - Used Capacity in GiB). For example, if the capacity limit is 80% or 80 GiB and the used capacity is 40 GiB, then the value for Capacity-to-Limit (GiB) is (80 GiB - 40 GiB or 80% - 50%) which is 30% or 40 GiB.

See these related values for more information Capacity Limit (%) and Adjusted Used Capacity (%).

This metric is not available for all storage systems, such as FlashSystem A9000, FlashSystem A9000R, and Dell EMC VMAX.

Compression Savings (%)

The estimated amount and percentage of capacity that is saved by using data compression. The percentage is calculated across all compressed volumes in the pool and does not include the capacity of non-compressed volumes.

For storage systems with drives that use inline data compression technology, the Compression Savings does not include the capacity savings that are achieved at the drive level.

The following formula is used to calculate the amount of storage space that is saved:

written capacity - compressed size

The following formula is used to calculate the percentage of capacity that is saved:

((written capacity - compressed size) ÷ written capacity) × 100

For example, the written capacity, which is the amount of data that is written to the volumes before compression, is 40 GiB. The compressed size, which reflects the size of compressed data that is written to disk, is just 10 GiB. Therefore, the compression savings percentage across all compressed volumes is 75%.

Availability: IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.

Deduplication Savings (%)

The estimated amount and percentage of capacity that is saved by using data deduplication. The percentage is calculated across all deduplicated volumes in the pool and does not include the capacity of volumes that are not deduplicated.

The following formula is used to calculate the amount of storage space that is saved:

written capacity - deduplicated size

The following formula is used to calculate the percentage of capacity that is saved:

((written capacity - deduplicated size) ÷ written capacity) × 100

For example, the written capacity, which is the amount of data that is written to the volumes before deduplication, is 40 GiB. The deduplicated size, which reflects the size of deduplicated data that is written to disk, is just 10 GB. Therefore, data deduplication reduced the size of the data that is written by 75%.

Availability: Storage systems that run IBM Spectrum Virtualize 8.1.3 or later.

Drive Compression Savings (%)

The amount and percentage of capacity that is saved with drives that use inline data compression technology. The percentage is calculated across all compressed drives in the pools.

The amount of storage space that is saved is the sum of drive compression savings.

The following formula is used to calculate the percentage of capacity that is saved:

((used written capacity - compressed size) ÷ used written capacity) × 100

Availability: Storage systems that contain IBM FlashCore Modules with hardware compression.

Enterprise HDD Available Capacity (GiB)

The amount of storage space that is available on the Enterprise hard disk drives that can be used by Easy Tier® for re-tiering the volume extents in the pool.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Enterprise HDD Capacity (GiB)

The total amount of storage space on the Enterprise hard disk drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Mapped Capacity (GiB)

(Previously known as Assigned Volume Space) The total amount of space in the volumes that is assigned to hosts. For Hitachi VSP non-thin provisioning pool space, this value is the sum of assigned regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.

Availability: All storage systems.

Nearline HDD Available Capacity (GiB)

The amount of storage space that is available on the Nearline hard disk drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Nearline HDD Capacity (GiB)

The total amount of storage space on the Nearline hard disk drives that can be used by Easy Tier for re-tiering the volume extents in the pool.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Overprovisioned Capacity (GiB)

(Previously known as Unallocatable Volume Space) The capacity that cannot be used by volumes because the physical capacity of the pool cannot meet the demands for provisioned capacity. The following formula is used to calculate this value:

[Provisioned Capacity - Capacity]

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool by creating volumes with more provisioned capacity than can be physically allocated in the pool. This value represents the amount of volume capacity that cannot be allocated based on the current capacity of the pool. For Hitachi VSP non-thin provisioned pool capacity, this value is always zero.

Availability: All storage systems, except FlashSystem A9000, FlashSystem A9000R, XIV, and IBM Spectrum Accelerate.

Provisioned Capacity (%)

(Previously known as Virtual Allocation) The percentage of the physical capacity that is committed to the provisioned capacity of the volumes in the pool. If the value exceeds 100%, the physical capacity doesn't meet the demands for provisioned capacity. The following formula is used to calculate this value:

[(Provisioned Capacity ÷ Capacity) × 100]

This value is available for all pools.

For Hitachi VSP non-thin provisioned pool space, the following formula is used to calculate this value:

[(Used Capacity ÷ Capacity) × 100]

Example: If the provisioned capacity percentage is 200% for a total storage pool size of 15 GiB, then the provisioned capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much capacity is committed than is physically contained in the pool. If the provisioned capacity percentage is 100% for the same pool, then the provisioned capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already used by volumes.

A provisioned capacity percentage that is higher than 100% is considered aggressive because insufficient physical capacity is available in the pool to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Shortfall (%) to estimate how critical the shortage of capacity is for a pool.

You can hover the mouse pointer over the percentage bar to view values for the provisioned capacity and capacity.

Availability: All storage systems.

Provisioned Capacity (GiB)

(Previously known as Total Volume Capacity) The total amount of storage capacity that can be made available to the standard- and thin-provisioned volumes in the pool. If the pool is a parent pool, it also includes the storage capacity that can be made available to the volumes in the child pools. For Hitachi VSP non-thin provisioned pool capacity, this value is the sum of the capacity of regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.

Availability: All storage systems.

Repository Capacity (GiB)

The total storage capacity of the repository for Track Space-Efficient (TSE) thin-provisioning.
Availability: DS8000 thin-provisioned pools.

Reserved Volume Capacity

(Previously known as Unused Space) The amount of pool capacity that is reserved but has not been used yet to store data on the thin-provisioned volume.
Availability: Resources that run IBM Spectrum Virtualize.

Safeguarded Capacity (GiB)

The total amount of capacity that is used to store volume backups that are created by the Safeguarded Copy feature in DS8000.

SCM Available Capacity (GiB)

The available capacity on Storage Class Memory (SCM) drives in the pool. Easy Tier can use these drives to retire the volume extents in the pool.

Availability: IBM Spectrum Virtualize systems, such as FlashSystem 9100, FlashSystem 7200, and Storwize family storage systems that are configured with block storage.

SCM Capacity (GiB)

The total capacity on Storage Class Memory (SCM) drives in the pool. Easy Tier can use these drives to retire the volume extents in the pool.

Availability: IBM Spectrum Virtualize systems, such as FlashSystem 9100, FlashSystem 7200, and Storwize family storage systems that are configured with block storage.

Shortfall (%)

The difference between the remaining unused volume capacity and the available capacity of the associated pool, expressed as a percentage of the remaining unused volume capacity. The shortfall represents the relative risk of running out of space for overallocated thin-provisioned volumes. If the pool has sufficient available capacity to satisfy the remaining unused volume capacity, no shortfall exists. As the remaining unused volume capacity grows, or as the available pool capacity decreases, the shortfall increases and the risk of running out of space becomes higher. If the available capacity of the pool is exhausted, the shortfall is 100% and any volumes that are not yet fully allocated have run out of space.


If the pool isn't thin-provisioned, the shortfall percentage equals zero. If shortfall percentage isn't calculated for the storage system, the field is left blank.

The following formula is used to calculate this value:

Overprovisioned Capacity ÷ Committed but Unused Capacity

You can use this percentage to determine when the amount of over-committed space in a pool is at a critically high level. Specifically, if the physical space in a pool is less than the committed provisioned capacity, then the pool does not have enough space to fulfill the commitment to provisioned capacity. This value represents the percentage of the committed provisioned capacity that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

Simplifying storage, deploying new applications, and controlling costs with IBM Spectrum Storage

 Simplifying storage, deploying new applications, and controlling costs with IBM Spectrum Storage

Example: The remaining physical capacity of a pool is 70 GiB, but 150 GiB of provisioned capacity was committed to thin-provisioned volumes. If the volumes are using 50 GiB, then 100 GiB is still committed to the volumes (150 GiB – 50 GiB) with a shortfall of 30 GiB (70 GiB remaining pool space – 100 GiB remaining commitment of volume space to the volumes).

Because the volumes are overcommitted by 30 GiB based on the available capacity in the pool, the shortfall is 30% when the following calculation is used:

**[(100 GiB unused volume capacity – 70 GiB remaining pool capacity)
÷ 100 GiB unused volume capacity] × 100**

Availability: DS8000, Hitachi Virtual Storage Platform, and storage systems that run IBM Spectrum Virtualize.

For FlashSystem A9000 and FlashSystem A9000R, this value is not available.

Soft Capacity (GiB)

The amount of virtual storage space that is configured for the pool.

Availability: XIV systems and IBM Spectrum Accelerate storage systems.

Tier 0 Flash Available Capacity (GiB)

The amount of storage space that is available on the Tier 0 flash solid-state drives that can be used by Easy Tier for retiring the volume extents in the pool.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier 0 Flash Capacity (GiB)

The total amount of storage space on the Tier 0 flash solid-state drives that can be used by Easy Tier for retiring the volume extents in the pool.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier 1 Flash Available Capacity (GiB)

The amount of storage space that is available on the Tier 1 flash, read-intensive solid-state drives that can be used by Easy Tier for retiring the volume extents in the pool.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier 1 Flash Capacity (GiB)

The total amount of storage space on the Tier 1 flash, read-intensive solid-state drives that can be used by Easy Tier for retiring the volume extents in the pool.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier 2 Flash Available Capacity (GiB)

The available capacity on Tier 2 flash, high-capacity drives in the pool. Easy Tier can use these drives to retire the volume extents in the pool.

Availability: DS8000 storage systems.

Tier 2 Flash Capacity (GiB)

The total capacity on Tier 2 flash, high-capacity drives in the pool. Easy Tier can use these drives to retire the volume extents in the pool.

Availability: DS8000 storage systems.

Total Capacity Savings (%)

The estimated amount and percentage of capacity that is saved by using data deduplication, pool compression, thin provisioning, and drive compression, across all volumes in the pool.

The following formula is used to calculate the amount of storage space that is saved:

Provisioned Capacity – Used Capacity

The following formula is used to calculate the percentage of capacity that is saved:

$$((\text{Provisioned Capacity} - \text{Used Capacity}) \div \text{Provisioned Capacity}) \times 100$$

Availability: FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.

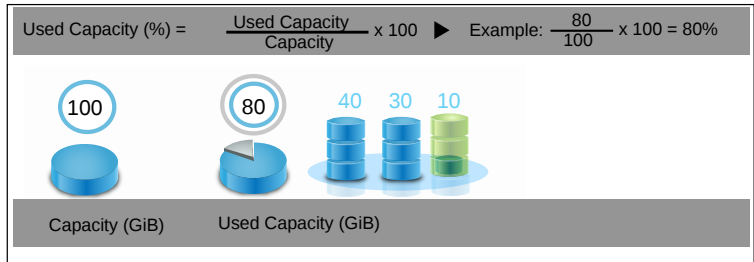
Unmapped Capacity (GiB)

(Previously known as Unassigned Volume Space) The total amount of space in the volumes that are not assigned to hosts. For Hitachi VSP non-thin provisioning pool space, this value is the sum of unassigned regular host-accessible volumes. Volumes that are used for thin-provisioning (pool volumes) are not included.

Availability: All storage systems.

Used Capacity (%)

(Previously known as Physical Allocation)



The percentage of physical capacity that is used by the volumes in the pool, including the volumes in child pools. This value is always less than or equal to 100% because you cannot allocate more physical space than is available in a pool. Check the value for used capacity to see:

- Whether the physical capacity of the pool is fully allocated. That is, the value for used capacity is 100%.
- Whether you have sufficient capacity to provision new volumes with storage
- Whether you have sufficient capacity to allocate to the compressed and thin-provisioned volumes in the pool

Availability: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Used Capacity (GiB)

(Previously known as Allocated Space) The amount of physical capacity that is used by the volumes in the pool. If the pool is a parent pool, the amount of space that is used by the volumes in the child pools is also included.

The capacity that is used by thin-provisioned volumes is less than their provisioned capacity, which is shown in the Provisioned Capacity column. If a pool does not contain thin-provisioned volumes, this value is the same as Provisioned Capacity.

Availability: All storage systems, except FlashSystem A9000 and FlashSystem A9000R.

Used Written Capacity (%)

(Previously known as Effective Used Capacity) For devices with inline hardware compression, the effective used capacity percentage is the percentage of capacity that is provisioned to the standard-provisioned volumes and the thin-provisioned volumes, given the drive compression savings.

Used Written Capacity (GiB)

(Previously known as Effective Used Capacity) The amount of capacity that is written to the volumes in a pool before inline disk compression is applied. If a pool is not compressed, this value is the same as Used Capacity.

Written Capacity Limit (GiB)

(Previously known as Effective Capacity) The maximum of amount of capacity that can be written to a pool before inline-disk compression is applied. If a pool is not compressed, this value is the same as Capacity.

Volume capacity metrics

You use the capacity chart to detect capacity shortages for the following types of volumes:

- Space-efficient volumes such as compressed volumes and thin-provisioned volumes
- Standard-provisioned volumes that use Easy Tier to re-tier volume extents

You can review the capacity usage by space-efficient volumes to detect capacity shortfalls. You can also review the capacity usage of volumes that use Easy Tier to distribute volume extents across Enterprise HDD, Nearline HDD, and SSD drives.

To detect capacity shortages and investigate capacity usage trends, you can add the following metrics to the chart for volumes:

Capacity (GiB)

The capacity of the compressed or the thin-provisioned volume, which comprises the sum of the used and available capacity. For thin-provisioned volumes in XIV systems pools or IBM Spectrum Accelerate pools, capacity is the physical ("hard") capacity of the volume.

Availability: All storage systems.

Compression Savings (%)

The estimated amount and percentage of capacity that is saved by using data compression.

The following formula is used to calculate the amount of storage space that is saved:

$$\text{written capacity} - \text{compressed size}$$

The following formula is used to calculate the percentage of capacity that is saved:

$$((\text{written capacity} - \text{compressed size}) \div \text{written capacity}) \times 100$$

Availability: FlashSystem A9000 and FlashSystem A9000R, IBM Spectrum Accelerate, XIV storage systems with firmware version 11.6 or later, and resources that run IBM Spectrum Virtualize.

Exception: For compressed volumes that are also deduplicated, on storage systems that run IBM Spectrum Virtualize, this column is blank.

Enterprise HDD Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Enterprise hard disk drives.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Nearline HDD Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Nearline hard disk drives.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Safeguarded Capacity (GiB)

The amount of capacity that is used to store volume backups that are created by the Safeguarded Copy feature in DS8000.

SCM Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Storage Class Memory (SCM) drives.

Availability: IBM Spectrum Virtualize systems, such as FlashSystem 9100, FlashSystem 7200, and Storwize family storage systems that are configured with block storage.

Tier 0 Flash Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Tier 0 flash drives.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier 1 Flash Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Tier 1 flash, read-intensive drives.

Availability: DS8000 and storage systems that run IBM Spectrum Virtualize.

Tier 2 Flash Capacity (GiB)

The amount of volume capacity that Easy Tier has placed on Tier 2 flash, high-capacity drives.

Availability: DS8000 storage systems.

Used Capacity (GiB)

(Previously known as Allocated Space) The amount of space that is used by the compressed, thin-provisioned, or the Easy Tier volume. Typically, the space that is used by the compressed or thin-provisioned volume is less than the capacity of the volume. For Easy Tier volumes, used capacity is the capacity that is used by the volume's extents on the Enterprise HDD, Nearline HDD, or SSD drives.

Availability: All storage systems.

Written Capacity (GiB)

(Previously known as Written Space) The amount of data that is written from the assigned hosts to the volume before compression or data deduplication are used to reduce the size of the data. For example, the written capacity for a volume is 40 GiB. After compression, the volume used space, which reflects the size of compressed data that is written to disk, is just 10 GiB.

Limitations and known issues

Zero values are displayed for some metrics on storage systems that run IBM Spectrum Virtualize

Storage systems that run IBM Spectrum Virtualize with firmware 8.2 or earlier might show zero values for some capacity values. This issue occurs because the capacity values for these storage systems were changed to rely on physical capacity values.

The issue is resolved in newer versions of firmware, starting with the following versions:

- 8.1.3.6
- 8.2.1.4

For more information, see [APAR HU01916: The GUI Dashboard and the CLI lssystem command report physical capacity incorrectly](#).

- [Key capacity concepts](#)

Understand the terms that are used to measure and plan capacity.

- [Key storage values for pools](#)

The metadata that is collected about storage systems is used to calculate key capacity values for pools. Use these values to monitor storage usage and to detect and identify capacity issues.

Key capacity concepts

Understand the terms that are used to measure and plan capacity.

Scenarios are used to explore the following capacity concepts for block storage systems:

- [Standard provisioning](#)
- [Thin-provisioning](#)
 - [Thin-provisioning savings](#)
 - [Data reduction savings](#)
 - [Total capacity savings](#)
 - [Effective capacity](#)

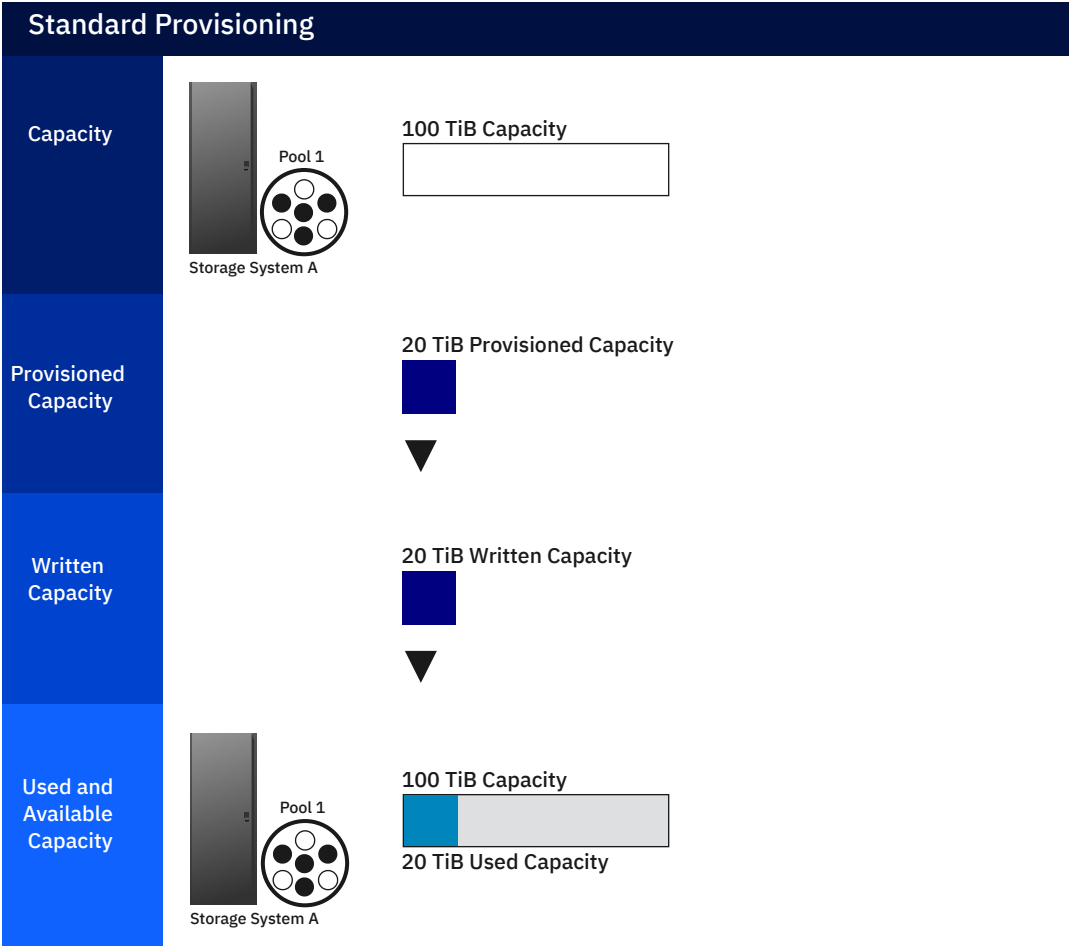
Tip: To review the capacity terms and definitions that are used, click [Definitions for the key capacity terms and concepts](#).

Standard provisioning

You add Storage System A with a capacity of 100 TiB to IBM Spectrum® Control for monitoring.

Then, you add a standard-provisioned volume with a capacity of 20 TiB.

You create standard-provisioned volumes when you want to dedicate the usable capacity of the storage system to the device that writes to it. From the perspective of the pool and storage system, the capacity is used and is no longer available to the storage system and pool.



The capacity values that are reported by the storage system for standard-provisioned volumes are as follows:

- Used Capacity (%) is reported as fully allocated, and Used Capacity (GiB) has the same value as the provisioned capacity of the volume.
- Available Capacity (%) is reported as 00.00.

The capacity of the data that is written to standard-provisioned volumes is not measured by the storage system. However, the written capacity of standard-provisioned volumes is included in the calculation of the total written capacity of the volumes. Because the data that is written to standard-provisioned volumes isn't reduced when data is stored, the written capacity has the same value as the provisioned capacity of the volume.

Figure 1. Overview chart: Written capacity of standard and thin-provisioned volumes



Tip: From the Storage menu, click Block Storage Systems. Double-click a storage system and check the Provisioned Capacity chart.


Thin-provisioning

Unlike standard-provisioned volumes, the capacity of thin-provisioned volumes is not dedicated to the device that writes to it. Devices get the capacity when they write data to the thin-provisioned volumes. Depending on the data reduction techniques that are supported by the storage system, the data that is written to the thin-provisioned volumes can be reduced before it is stored on the volumes.

Thin-provisioning savings

Thin-provisioning savings are the total amount of capacity that is saved in a pool, system, or volume by using capacity when needed as a result of write operations. The capacity that is saved is the difference between the provisioned capacity and the written capacity.

In this scenario, you add four thin-provisioned volumes with a total capacity of 200 TiB to Storage System B, which has a capacity of 100 TiB.



Thin-provisioning savings are the difference between the 200 TiB total provisioned to the thin-provisioned volumes and the 150 TiB of data that is written to the volumes, which is 50 TiB.

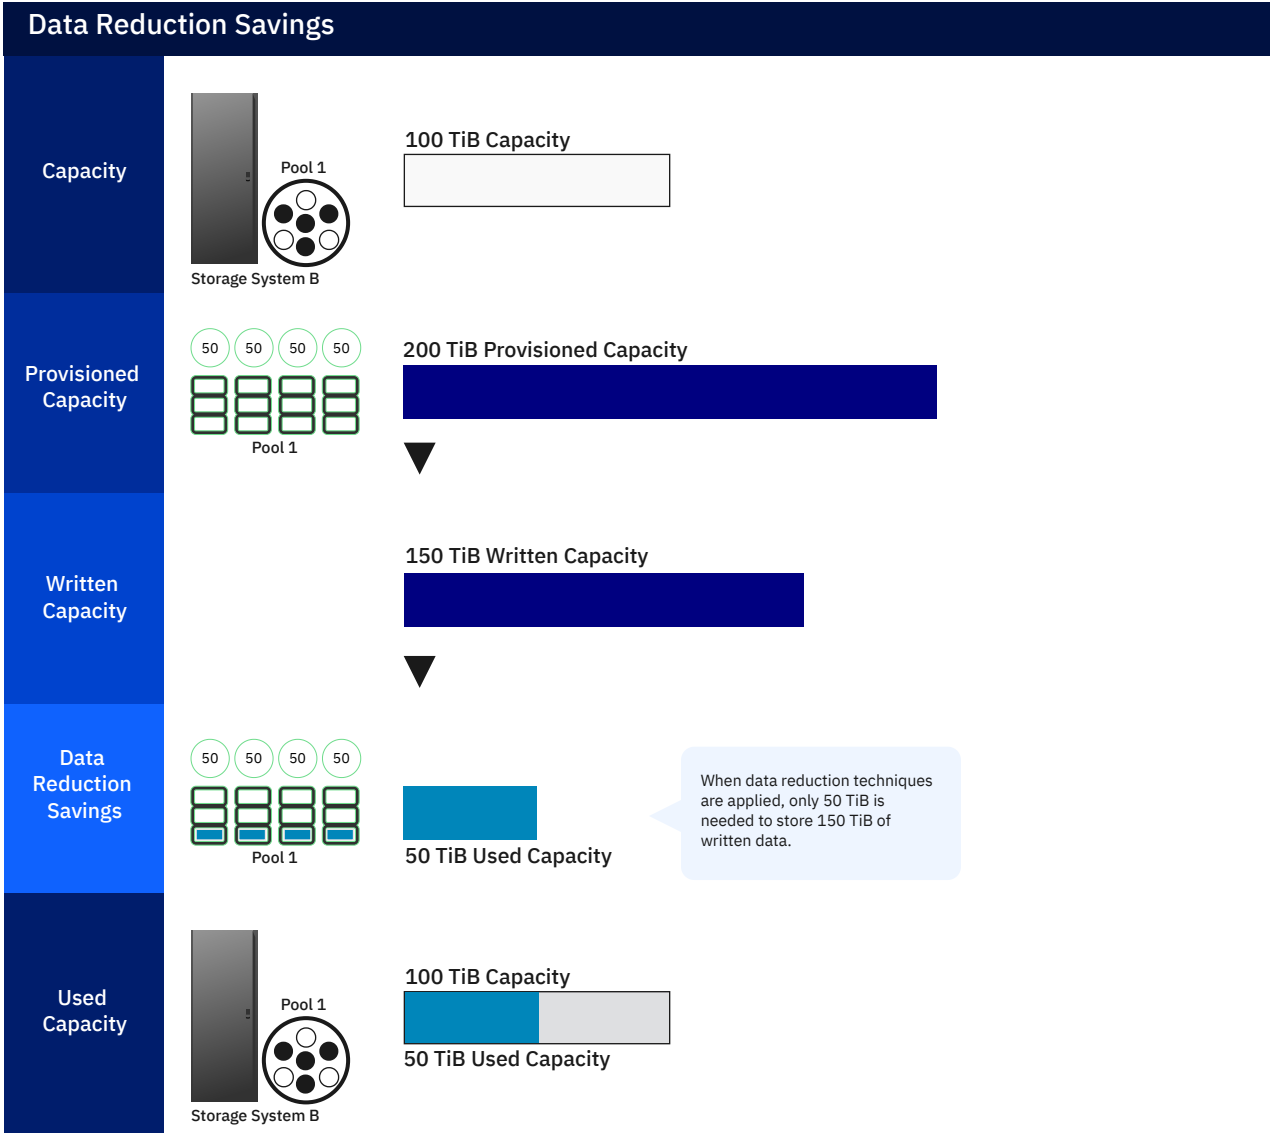
Measuring the capacity data reduction savings of thin-provisioned volumes  Measuring the capacity data reduction savings of thin-provisioned volumes

(To enlarge, right-click the video and select the full screen or new tab option.)

Data reduction savings

Data reduction savings are the total amount of usable capacity that is saved in a pool, system, or volume through the application of data reduction algorithms on the written data, such as compression and deduplication. This saved capacity is the difference between the written capacity and the used capacity.


Before data is stored on thin-provisioned volumes, it can be reduced. For example, if your storage system supports compression and deduplication, these data reduction techniques are applied to reduce the amount of capacity that is needed to store the data.



In this scenario, the data reduction savings are the difference between the written capacity (150 TiB), and the used capacity(50 TiB), which is 100 TiB or 67%.

Total capacity savings

Total capacity savings are the amount of capacity that is saved in a pool, system, or volume through thin-provisioning and data reduction techniques.



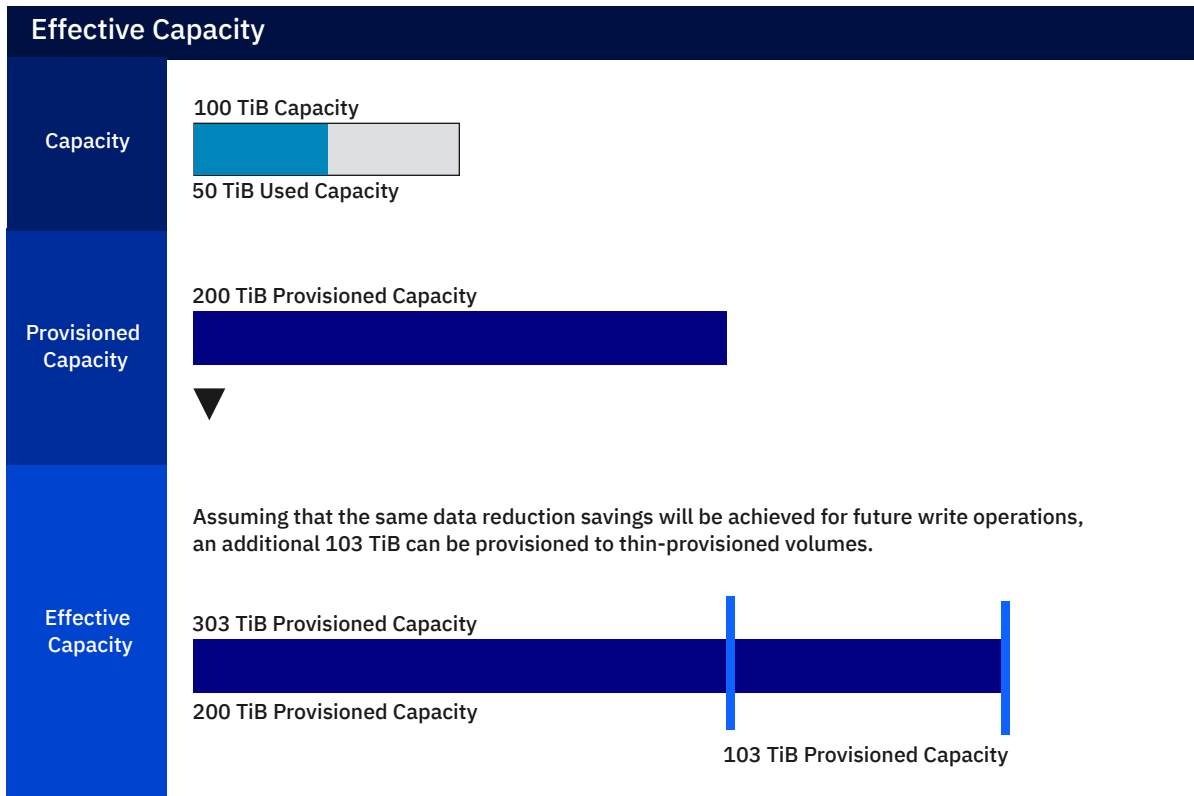
In the scenario illustrated above, the thin-provisioning savings are (200 TiB - 150 TiB), which is 50 TiB. The data reductions savings are (150 TiB - 50 TiB), which is 100 TiB. So, the total capacity savings are 150 TiB.

Effective capacity

Effective capacity is the amount of provisioned capacity that can be created in a system or pool without running out of capacity given the current data reduction savings that are being achieved. To calculate effective capacity, you divide the capacity of the storage system or pool by the data reduction savings percentage.

Note: In some storage systems, restrictions in the system determine the maximum provisioned capacity that is allowed in a pool or system. In these cases, the effective capacity cannot exceed this limit.

You can use effective capacity to estimate how much capacity you can provision to your thin-provisioned volumes without running out of usable capacity.



In the scenario illustrated above, the storage system has a capacity of 100 TiB and a current data reduction savings of 66%. So, the effective capacity is $100 / (1.00 - 0.67)$, which is 303 TiB. Currently, the provisioned capacity of the volumes is 200 TiB. Assuming that the same data reduction savings are achieved for subsequent write operations, more thin-provisioned volumes can be created with a total capacity of 303 TiB.

Definitions for the key capacity terms and concepts

These are capacity terms and concepts that are used:

Available capacity

The amount of usable capacity that is not yet used in a system, pool, array, or MDisk.

Capacity

The amount of usable capacity that is available for storing data on a system, pool, array, or managed disk after formatting and RAID techniques are applied.

Compression

A function that removes repetitive characters, spaces, strings of characters, or binary data from the data being processed and replaces characters with control characters. Compression reduces the amount of storage space that is required for data.

Data deduplication

A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage media. Other instances of the same data are replaced with a pointer to the retained instance.

Data reduction

A set of techniques that can be used to reduce the amount of usable capacity that is required to store data. Examples of data reduction include data deduplication and compression.

Data reduction savings

The total amount of usable capacity that is saved in a pool, system, or volume through the application of data reduction algorithms on the written data, such as compression and deduplication. This saved capacity is the difference between the written capacity and the used capacity.

Effective capacity

The amount of total provisioned capacity that can be created in a system or pool without running out of usable capacity given the current data reduction savings that are being achieved. This capacity equals the usable capacity that is divided by the data reduction savings percentage. In some storage systems, restrictions in the system determine the maximum provisioned capacity that is allowed in a pool or system. In those cases, the effective capacity cannot exceed this limit.

Provisioned capacity

The total capacity of all volumes in a pool or system.

Standard-provisioned volumes

Unlike thin-provisioned volumes, which use capacity when it is needed, the capacity that is provisioned to standard-provisioned or thick volumes is no longer available to the pool. Standard-provisioned volumes are fully allocated, and their capacity is reported as used because their capacity is no longer available to the pool and storage system.

Thin-provisioning

The ability to defer capacity allocation on a storage resource until data is actually written to it.

Thin-provisioning savings

The total amount of usable capacity that is saved in a pool, system, or volume by using usable capacity when needed as a result of write operations. The capacity that is saved is the difference between the total provisioned minus the written capacity.

Total capacity savings

The total amount of capacity that is saved in a pool, system, or volume through thin-provisioning and data reduction techniques. This saved capacity is the difference between the used capacity and the provisioned capacity.

Used capacity

The amount of capacity that is taken up by data in a system, pool, array, or MDisk after data reduction techniques have been applied.


Written capacity

The amount of usable capacity that would have been used to store written data in a pool or system if data reduction was not applied.

Key storage values for pools

The metadata that is collected about storage systems is used to calculate key capacity values for pools. Use these values to monitor storage usage and to detect and identify capacity issues.

The following diagram shows the relationships between the key storage values for pools:



The following table describes the values on the diagram and how they are calculated:

Storage value	Formula	Description
Capacity	A	The total amount of physical and usable capacity in a pool.
Used Capacity	B	The amount of physical capacity that is used by the volumes in a pool, including the volumes in child pools.
Available Volume Capacity	$PC - B$	The amount of the Provisioned Capacity (PC) in the pool that is not used.
Used Capacity %	$B \div A \times 100$	The percentage of physical capacity that is used by the volumes in a pool, including the volumes in child pools.
Provisioned Capacity %	$PC \div A \times 100$	The percentage of the physical capacity that is committed or reserved for the provisioned capacity of the volumes in the pool.
Reserved Volume Capacity	$B - C$	The amount of pool capacity that is reserved but has not been used yet to store data on the thin-provisioned volume.
Provisioned Capacity (PC)	Sum of volume capacities	The total storage capacity on all the volumes in a pool, which includes thin-provisioned and standard-provisioned volumes.
Overprovisioned Capacity %	$E = PC - A$	The amount of capacity by which the Provisioned Capacity exceeds the physical capacity of a pool.
Shortfall %	$((E \div (D + E)) \times 100$ $E =$ Overprovisioned Capacity %	The percentage of the remaining capacity in a pool that is not available to be used. The higher the percentage, the more critical the shortfall of pool capacity.
Reserved Capacity	F	The amount of unused capacity in a pool that is reserved by provisioning and optimization tasks. Pool capacity is reserved when a provisioning or optimization task is created, and used when the task is run.

Capacity metrics for file storage systems

To review trends in capacity and space usage for file storage systems, you add metrics to capacity charts. You use the charts for filesets, file systems, and file system pools to detect capacity shortages and space usage trends.

Alphabetical lists of the capacity and space usage metrics that you can add to charts are provided in the following sections:

- [Capacity metrics for filesets](#)
- [Capacity metrics for file systems](#)
- [Capacity metrics for file system pools](#)

Capacity metrics for filesets

If sufficient data is collected about the filesets in your file storage systems, you can view charts that compare the used capacity of the filesets. The following metric is displayed in the capacity chart for filesets:

Used Capacity (GiB)

The amount of storage capacity that is used by the fileset. Used capacity is not provided for filesets that are cache targets.

Availability: All storage systems.

Capacity metrics for file systems

To detect capacity shortages and investigate capacity usage trends, you can add the following metrics to the capacity chart for file systems:

Available Inodes

The number of inodes that are available in a file system.

Available Capacity (GiB)

The amount of storage capacity that is available (not allocated) on a file system.

Capacity (%)

The percentage of the total storage space on the file system that is used by files and directories.

Available for: All storage systems.

Capacity (GiB)

(Previously known as Total Capacity) The capacity on the file system.

Available for: All storage systems.

Maximum Inodes

The total number of inodes that the file system can contain. This value consists of the available inodes and the used inodes for the file system.
Available for: All storage systems.

Physical Capacity (GiB)

The raw capacity of the partition where a file system resides.

Used Inodes

The number of used inodes on a file system. An inode is the internal structure that describes the individual files or directories in the file system metadata. An inode contains the node, type, owner, and location of a file or directory.

Availability: All storage systems.

Used Inodes (%)

The percentage of inodes that are already used on the file system. Each time that you create a file or directory on the file system, an inode is allocated to the file or directory.

Available for: All storage systems.

Internal Used Capacity (GiB)

The amount of storage space that is unavailable (allocated) on a file system. For IBM Spectrum Scale, this value does not include the space that is used by migrated data on external pools.

Capacity metrics for file system pools

If sufficient data is collected about the file system pools in your data center, you can view charts that compare the capacity of the pools with the space that is allocated to the pools and the space that is still available in the pools.

The internal resource for IBM Spectrum Scale storage systems is called Pools. For all other file storage systems, the internal resource is called File System Pools.

To detect capacity shortages and investigate space usage trends, you can add the following metrics to the capacity chart for file system pools:

Available Capacity (GiB)

The amount of unallocated storage capacity in the pool. Available capacity usually comprises the space that can be used for storage. However, if the pool is not formatted, the amount of overhead capacity might also be included in the calculation.

Availability: All storage systems.

Capacity (%)

The percentage of space in the file system pool that is being used.

Available for: All storage systems.

Capacity (GiB)

(Previously known as Total Capacity) The capacity in the storage pool.

Available for: All storage systems.

Used Capacity (GiB)

The amount of space in the file system pool that is being used.

Available for: All storage systems.

Capacity metrics for object storage systems

To review trends in capacity and space usage for object storage systems, you add metrics to capacity charts. Use the charts for containers to detect capacity shortages and space usage trends for the containers in your object storage systems.

Capacity metrics for containers

If sufficient data is collected about the containers in your object storage systems, you can view charts that compare the used capacity of the containers. You can add the following metrics to the capacity chart for containers:

Available Objects

The number of objects that you can add to the container on the object storage system.

Available for: IBM Spectrum Scale storage systems.

Available Capacity (GiB)

The amount of file system capacity that can be used to store object data for the container.

Available for: IBM Spectrum Scale storage systems.

Capacity Quota (%)

The percentage of the capacity quota for the container that was used when the storage system was last probed. The quota limits the amount of file system capacity that can be used by the container. This column contains a value only if a capacity quota is defined for the container on the OpenStack Swift object storage system. If the used capacity value for the container exceeds the capacity quota value, the percentage exceeds 100%. The used capacity can exceed the capacity quota if the quota is set after object files are uploaded to the file system or if an object file without size information is uploaded to the file system.

Available for: IBM Spectrum Scale storage systems.

Capacity Quota (GiB)

The amount of file system capacity that can be used to store object data for the container. This column contains a value only if a capacity quota is defined for the container on the OpenStack Swift object storage system.

Available for: IBM Spectrum Scale storage systems.

Objects

The number of objects in the container on the object storage system.

An object stores data content, such as files, videos, images, virtual machine snapshots, and other unstructured data. In IBM Spectrum Scale, objects are stored as files on the GPFS file system.

Available for: IBM Spectrum Scale storage systems.

Objects Quota

The number of objects that can be stored in the container. This column contains a value only if an objects quota is defined for the container on the OpenStack Swift object storage system.

Available for: IBM Spectrum Scale storage systems.

Objects Quota (%)

The percentage of the objects quota for the container that was used when the storage system was last probed. The quota limits the number of objects that can be stored in the container. This column contains a value only if an objects quota is defined for the container on the OpenStack Swift object storage system. If the number of objects that are used exceeds the objects quota, the percentage exceeds 100%. The number can exceed the quota if the quota is configured after the object files are stored on the file system.

Available for: IBM Spectrum Scale storage systems.

Used Capacity (GiB)

The amount of file system capacity that is used by the objects in the container. If a probe is run immediately after objects are added to a container, the used capacity value might not reflect the updates until after the next scheduled probe is run.

Available for: IBM Spectrum Scale storage systems.

Capacity metrics for tiers

Review trends in the capacity and space usage for tiers. You use the charts to detect capacity shortages and space usage for the tiers in your storage environment.

Based on historical space usage, you can plan the capacity growth for tiers.

You can view values for the following metrics in the tier charts:

Capacity (GiB)

The total amount of storage space in the tier, which comprises the total storage space of the pools in the tier.

Used Capacity (GiB)

The capacity that is used by pools in the tier.

Monitoring and administering applications, departments, and general groups

You create applications to monitor the storage capacity, space usage, and performance of applications, and you create departments to monitor the space usage of the applications in the department. You can structure your applications and departments hierarchically to match the structure of your business organization.

You can map the application to the storage resources that the application uses to do capacity trending, health monitoring, and performance troubleshooting tasks. For example, a department might use 15 applications and be part of another five departments. Also a department might share storage resources with another department, subdepartment or application even if they do not belong in the same business hierarchy. Another example, might be if you are contacted because an application within a department is experiencing a performance issue. You carry out tasks like capacity reporting or trending on the behalf of the application.

Create general groups to quickly view information about storage resources that have common characteristics. For example, you might group the subset of ports on a SAN Volume Controller that are used for inter-node communication, or all the storage systems with lease agreements that end in the current year, or the storage systems that are used by a critical business application.

Application administration

Go to the pages for creating and administering the applications that you want to monitor.

Table 1. Create and administer applications and subcomponents

Actions	Navigation
Create applications and assign storage resources directly to the application.	<ol style="list-style-type: none">1. In the menu bar, go to Groups > Applications.2. On the Applications page, click Create Application.3. On the Select a Method page, click Assign storage resources to the application.
Create subcomponents that can have its own storage resources.	<ol style="list-style-type: none">1. In the menu bar, go to Groups > Applications.2. On the Applications page, click Create Application.3. On the Select a Method page, click Create a subcomponent that can have its own set of storage resources.
Create a filter to use in assigning resources.	<ol style="list-style-type: none">1. In the menu bar, go to Groups > Applications.2. On the Applications page, right-click a row in the Applications list and click View Details.3. In the Members section, click Filters.4. On the Filters page, click Create Filter.
Add applications as subcomponents to an application.	<ol style="list-style-type: none">1. In the menu bar, go to Groups > Applications.2. Right-click the applications that you want to add as subcomponents, and then click Add to.. > Application.3. On the Add to Application page, select the applications and click Save.
Add applications to departments.	<ol style="list-style-type: none">1. In the menu bar, go to Groups > Applications.2. On the Applications page, select one or more applications, right-click and then click Add to.. > Department.3. On the Add to Department page, select the applications that you want to add as members and click Save.
Add applications to subdepartments.	<ol style="list-style-type: none">1. In the menu bar, go to Groups > Departments.2. Right-click the department, and then click View Details.3. In the General section, click Subdepartments.4. Right-click the subdepartment, and click View Details.5. In the General section, click Applications.6. Click Create Application.
Remove applications and subcomponents that you do not want to monitor.	<ol style="list-style-type: none">1. In the menu bar, go to Groups > Applications.2. On the Applications page, select one or more applications, and click Delete.3. If you select Delete subcomponents, the subcomponents that belong only to that application are removed.

Actions	Navigation
Remove subcomponents within direct context of an application that you do not want to monitor.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups, > Applications. 2. Right-click the application and click View Details. 3. In the General section, click Subcomponents. 4. Right-click the subcomponent, and click Remove from Application.
Remove selected resources from applications or subcomponents	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups, > Applications. 2. Right-click the application and click View Details. 3. In the Related Resources section, click an eligible resource. 4. Optional: In the Members section, click the storage resource that is part of the application. 5. On the resource page, right-click the resource, and click Remove from Application.
Troubleshoot applications	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups, > Applications . 2. Right-click the application, and then click View Details. <p>Investigate the performance of the application and the storage resources that are associated with the application by viewing the charts on the Overview page. You can investigate the performance of related resources by clicking the resource and clicking the Performance tab.</p>
Troubleshoot subcomponents	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups, > Applications . 2. Right-click the application, and then click View Details. 3. In the General section, click Subcomponents. 4. Right-click the subcomponent, and click View Details. <p>Investigate the performance of the storage resources that are associated with the subcomponent by viewing the charts on the Overview page. You can investigate the performance of related resources by clicking the resource and clicking the Performance tab.</p>

Department administration

Go to the pages for creating and administering the departments that you want to monitor.

Table 2. Create and administer departments and subdepartments

Actions	Navigation
Create departments and associate members (applications and departments) to create a hierarchal structure.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups, > Departments. 2. On the Departments page, click Create Department. 3. On the Add Members page, click Add one or more applications or Create subdepartment.
Add departments as subdepartments to a department.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups, > Departments. 2. On the Departments page, right-click the departments that you want to add as subdepartments, and then click Add to Department.
Add applications to departments.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups, > Departments. 2. Right-click the department, and then click View Details. 3. In the General section, click Applications. 4. Click Create Application.
Add applications to subdepartments.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups, > Departments. 2. Right-click the department, and then click View Details. 3. In the General section, click Subdepartments. 4. Right-click the subdepartment, and click View Details. 5. In the General section, click Applications. 6. Click Create Application.
Remove departments, subdepartments, and applications, that are members of the department.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups, > Departments. 2. On the Departments page, right-click one or more departments and click Delete. 3. If you select Delete subdepartments and applications, the subdepartments and applications that belong only to the department, that is being removed, are removed.
Remove applications from a department or subdepartment.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups, > Departments. 2. Right-click the department, and then click View Details. 3. In the General section, click Subdepartments. 4. Right-click the subdepartment, and click View Details. 5. In the General section, click Applications. 6. Select and right-click the applications, and then click Delete or Remove from Department.
Remove subdepartments within direct context of a department that you do not want to monitor.	<ol style="list-style-type: none"> 1. In the menu bar, go to Groups, > Departments. 2. Right-click the department, and then click View Details. 3. In the General section, click Subdepartments. 4. Right-click the subdepartments that you want to remove, and then click Remove from Department.

General group administration

Administer the general groups that you use to organize storage resources with common characteristics.

Table 3. Administering general groups

Actions	Navigation
---------	------------

Actions	Navigation
Create a general group and add resources.	<ol style="list-style-type: none"> 1. In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add file storage systems, click Storage > File Storage Systems. 2. Right-click one or more resources and click Add to General Group. 3. Click Add to new group, specify a name and description for the new group, and click Save.
Add resources to general groups.	<ol style="list-style-type: none"> 1. In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add block storage systems, click Storage > Block Storage Systems. 2. Right-click one or more resources and click Add to General Group. 3. Take one of the following actions: <ol style="list-style-type: none"> a. To add the resources to one or more existing groups, click Add to existing groups, click the appropriate groups in the list, and click Save. b. To add the resources to a new group, click Add to new group, specify a name and description for the new group, and click Save.
Set which alert policy manages a general group Learn more	<ol style="list-style-type: none"> 1. Go to the details page for the general group for which you want to set the policy. In the menu bar, click Groups > General Groups. 2. Right-click one of the general groups, then click View Details. 3. Click Alerts in the General section. 4. Click Set Policy from the Policy Actions menu. <p>To set the alert policy for multiple general groups, go to Settings > Alert Policies. Double-click the policy, click the resources tab, then click Edit Resources.</p>
Create an alert policy from the alert definitions and notification settings in a general group	<ol style="list-style-type: none"> 1. Go to the details page for the general group from which you want to create the policy. In the menu bar, click Groups > General Groups. 2. Right-click one of the general groups, then click View Details. 3. Click Alerts Definitions in the General section. 4. Click Create Policy from the Policy Actions menu.
View or modify the alert policy that manages a general group Learn more about defining alerts and notification settings	<ol style="list-style-type: none"> 1. Go to the details page for the general group whose policy you want to view. In the menu bar, click Groups > General Groups. 2. Right-click one of the general groups, then click View Details. 3. Click Alerts in the General section. 4. Click View Policy from the Policy Actions menu.
Delete general groups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click one or more general groups and click Delete. 3. To delete subgroups, ensure that the Delete subgroups? check box is selected. <p>When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product.</p>
Remove resources from a general group.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click a general group and click View Details. 3. In the Members section, click the type of resource that you want to remove. For example, if you want to remove some of the volumes from the group, click Volumes. 4. Right-click the resources that you want to remove and click Remove from General Group. 5. Click OK to confirm that you want to remove the resources from the general group and from all of its subgroups.
Add general groups as subgroups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click one or more general groups and click Add to General Group. 3. Take one of the following actions: <ol style="list-style-type: none"> a. To add the groups as subgroups of existing groups, click Add to existing groups, click one or more groups in the list, and click Save. b. To add the groups as subgroups of a new group, click Add to new group, specify a name and description for the new group, and click Save.
Remove subgroups from a group hierarchy.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click a general group and click View Details. 3. Click Subgroups in the General section. 4. Right-click the general groups that you want to remove and click Remove from General Group. <p>When you remove a subgroup from its parent group, the subgroup is moved to the same level in the hierarchy as the parent group. The subgroup is still monitored by the product.</p>
Create a subgroup.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click the general group that you want to create a subgroup for and click View Details. 3. Click Subgroups in the General section. 4. Click Create Subgroup. 5. Specify a name and description for the new group. 6. Customize the icon for the group by clicking the existing icon and selecting another icon. 7. Click Create.

Actions	Navigation
Delete subgroups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups > General Groups. 2. Right-click the general group that contains the subgroups that you want to delete and click View Details. 3. Click Subgroups in the General section. 4. Right-click one or more subgroups and click Delete. 5. To delete further subgroups, ensure that the Delete subgroups? check box is selected. <p>When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product.</p>

- [Applications](#)
View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.
- [Departments](#)
View information about departments and structure the hierarchy to mirror your business organization. Monitor the detailed information about capacity and space usage that is collected for departments and subdepartments, and monitor the performance of the applications and the application subcomponents that are added to departments and subdepartments.
- [Application and department hierarchies](#)
The interactive breadcrumb on the Overview page provides a visual representation of the hierarchical relationships for applications and departments and enables you to navigate quickly between the elements in the hierarchy.
- [Exporting information about applications and departments](#)
You can save information about applications and departments to a PDF, CSV, or HTML file. Information that you can export from the GUI includes all the values that are being shown in the columns for a list of applications and departments.
- [General groups](#)
View and administer the general groups that are used to alert on logically related storage resources. You can define alerts for a general group to notify you about changes in the configuration, attributes, and performance of the resources in the group.

Related tasks

- [Monitoring capacity usage at different levels of a business hierarchy](#)
- [Comparing storage usage in each department](#)
- [Using applications and subcomponents to monitor capacity and space usage](#)
- [Viewing storage capacity and usage trends](#)

Applications

View information about storage area network resources that you can model by using applications in your business environment. The monitoring and management of applications enables you to perform capacity trending, health monitoring, and performance troubleshooting tasks.

An application can be part of a department and have subcomponents that can contain a five level deep hierarchy. Applications that are grouped together can range from large line of business systems to specialized software, in a department that runs on either client computers or servers. For example, an application might be an automated billing system in the Finance department, VMware running in the Information Technology department or an email marketing system that is part of the Marketing department.

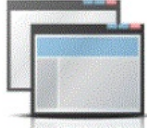
Another example might be an application named *Database* and the database contains logs and data, and a best practice is to place the logs and data on separate volumes. You might create subcomponents for the logs and data and these subcomponents become part of the *Database* application.

A storage administrator can use applications for planning purposes to diagnose how much storage a particular application is consuming in relation to the other applications in the business hierarchy. Another applicable usage for applications might be troubleshooting performance issues. The administrator checks the related resources that make up the application to see if there is a health status issue or can examine the performance of volumes that are used by the application.

The Applications page shows capacity information about the applications that are monitored and the actions that you can use to view and manage them. To see more detailed information, view properties and details for the application. If the application is associated with a department, information about that department is also shown.

The following information shows the resources and objects that can be added to the application.

Table 1. Resources and objects that can be added to the application

Object	Resources or objects that can be added to the application
 Application	<ul style="list-style-type: none"> • Applications • Data stores • Hypervisors • File Systems (Only file systems that are monitored through a Storage Resource Agent.) • Filesets • Servers • Shares • Vaults (IBM® Cloud Object Storage) • Virtual Machines (Only can be added with the CLI modifyappgroup command) • Volumes • Volume Groups

- [Creating an application hierarchy](#)
To monitor, the capacity, space usage, and performance of the storage resources that your applications use, create application models.
- [Adding resources manually to applications](#)
You can add resources directly from a resource page to your applications.
- [Adding resources with the command line interface](#)
You can use the **modifyappgroupviafile** command-line interface (CLI) command to automate your bulk assignment of eligible storage resources to applications. To

add a single resource to an application, use the **modifyappgroup** command.

- [Viewing information about applications](#)

You can view information about the application that you created as related resources, filters that are used to populate the application with resources, any subcomponents, capacity trending and performance information.

- [Removing applications and subcomponents](#)

You can remove applications and their subcomponents. The data that was collected to monitor the applications is also removed from the database repository.

- [Removing resources from applications](#)

Related resources within the context of an application can be removed from applications. The action is only available when viewing the related resources that belong to an application on the details page.

Related tasks

- [Defining alerts for applications](#)

Creating an application hierarchy

To monitor, the capacity, space usage, and performance of the storage resources that your applications use, create application models.

About this task

In this scenario, you want to monitor the capacity, space usage, and performance of a Db2® application.

The Db2 application has two subcomponents, which you want to monitor separately:

- logs
- data

To monitor the Db2 applications and the data and logs subcomponents, complete these tasks:

- Create an application that is named Db2.
- Create applications that are named data and logs
- Associate storage resources with the data and logs applications. For example, in this case, associate the volumes that the data application uses, type **data_vol***. For the logs application, type **logs_vol*** in the Volumes field of A selection of volumes that match a pattern filter. Then click Belonging to and click Storage System. Type the name of the storage system or the pattern that matches the name of the storage system.
- Add the data application and the logs application as subcomponents to the Db2 application.

Tip:

When you create complex application models for applications, it is not necessary to create a filter to associate storage resources with the application because the application inherits the storage resources that are associated with its subcomponents. However, if you want to monitor storage resources that are not associated with the subcomponents, you must create a filter that associates the additional storage resources that you want to monitor with the application.

Results

You can now view information about the capacity, space usage, and performance of the storage resources on the Overview page for the application and subcomponents.

What to do next

You can create departments and add applications to the departments.

1. [Creating applications](#)

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum® Control.

2. [Using filters to assign resources to applications](#)

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

3. [Adding subcomponents](#)

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

Related concepts

- [Application and department hierarchies](#)

Creating applications

To monitor storage capacity, troubleshoot performance, and view storage area network dependencies for applications, you can create applications for monitoring in IBM Spectrum Control.

About this task

Use the GUI to create applications, assign resources, and subcomponents for monitoring by IBM Spectrum Control.

To determine the storage capacity usage by the storage resources that the applications use, you create a data model of the applications for your storage environment.

Procedure

To create an application, complete these steps:

1. In the menu bar, go to Groups > Applications.
2. On the Applications page, click Create Application.
3. On the Create Application page, enter information about the application.

Note: You cannot have a duplicate application name on the same level in your application hierarchy. For example, you create an application named *Finance* and create a subcomponent named *Business*. You can create another application named *Business* as a top level application, but you cannot add the new *Business* application as a subcomponent of the *Finance* application.

You can customize the icon for your application by clicking the existing icon on the Create Application page and selecting another icon from the Customize Icon page.

4. Click Create.
5. Optional: On the Select a Method page, select Assign storage resources to the application that was created and follow the instructions to add the resources directly to the application.
6. Click Preview to see a list of filter matches before you finalize your filtering criteria, otherwise, click Save.
7. Optional: Select Create a subcomponent that can have its own set of storage resources and follow the instructions to create a subcomponent of the application.
8. Click Create.

Results

When the application creation is created, you can view status information, block and file capacity and additional data on the Applications page and in properties. If the View Performance action is available, you can view the aggregated information for all of the volumes in the application on the performance view. By default, the Total I/O Rate and the Total Data Rate are displayed in the performance chart. You can also select to display the read and write I/O rates, and the read and write data rates in the chart.

Note: The application performance feature isn't available to all IBM Spectrum Control users. Look for news from IBM® about when this feature will be available to you. For a detailed view of the application and related resources, on the Applications page, right-click the application and select View Details. To see the subcomponents for the application, click Subcomponent on the application details page.

Next topic: [Using filters to assign resources to applications](#)

Related tasks

- [Adding subcomponents](#)
- [Using filters to assign resources to applications](#)
- [Viewing information about applications](#)

Using filters to assign resources to applications

Instead of manually adding the resources that you want to assign to new or existing applications, you can create filters to assign resources based on name pattern matching.

About this task

When you create a filter to assign storage resources to applications, the resources that match the criteria that are defined in the filter are automatically added to the application or subcomponent (member application).

The use of wildcard characters enables you to customize and expand your search to add related resources. The asterisk (*) and the question mark (?) can be used as wildcard characters. The asterisk matches one or more characters. The question mark matches a single character. Use a comma-separated list to add multiple resources. As new resources are added or modified and they match the existing filter, they automatically become part of the application. For example, a new server is added, a volume is renamed or a share is mounted on a server.

If you edit a filter, the resources that no longer match the filter are removed from the application and the resources that now match the filter are added to the application.

If you remove the filter, all storage resources that match the filter criteria are removed from the application.

Filtering is extremely useful in storage environments where resources are predefined and named in a consistent format. Instead of adding each resource separately, you create filtering criteria that matches your naming format in your environment. You can use filtering when you create a new application or directly from the details page for an existing application.

You can only select one resource type filter at a time during the creation workflow, but you can create multiple filters for applications.

Tip: For any volume that you add to an application, the performance information for that volume, including historical information, is aggregated with the information for the existing volumes.

Note: The application performance feature isn't available to all IBM Spectrum® Control users. Look for news from IBM® about when this feature will be available to you. The filters are divided into the following categories:

All storage that belongs to a set of servers/hypervisors

You can add servers or hypervisors to the application or subcomponent (member application).

For example, in the Define Name Patterns section, enter ***blade**, in the Server/Hypervisor field and click Save. The filter adds all known servers or hypervisors that have a name that ends with **blade** to the application.

If you enter **server???.abc.com**, in the Server/Hypervisor field, the filter adds all servers, in the domain with names that begin with **server** followed by three characters and ends with **.abc.com** to the application. If you enter, **accounting*. ***, the filter adds all servers with names that begin with **accounting** regardless of the domain the system belongs to.

If you enter **cs-blade47.1ab.abc.com**, **cs-blade49.1ab.abc.com** or **cs-blade47***, **cs-blade49***, the filter adds multiple servers to the application.

A selection of volumes that match a pattern

You can add volumes of a certain name format, volumes that belong to a server or hypervisor, volumes from a specific storage system or volumes from a pool to the application or subcomponent. You can only select one resource option per filter creation.

For example, in the Define Name Patterns section, if you enter ***vo1**, in the Volumes field, mark the Belonging to box, click Server/Hypervisor, enter ***.lab.abc.com** and click Save. The filter adds all volumes that have a name that ends with **vo1** and are mapped to servers that have a name that ends with **.lab.abc.com** to the application.

If you enter ***vo1**, in the Volumes field, mark the Belonging to box, click Storage System, and enter **ds8000**. The filter adds all volumes that have a name that ends with **vo1** that belong to the **ds8000** storage system.

A selection of shares that match a pattern

You can add shares and exports that are mounted on a server or hypervisor, or shares and exports that are exported from a specific storage system to the application or subcomponent. You can select only one resource option per filter creation.

For example, in the Define Name Patterns section, if you enter ***share** in the Shares field, mark the Belonging to box, click Server/Hypervisor, enter ***.lab.abc.com** and click Save. The filter adds shares that have a name that ends with **share** that are mounted on servers that have a name that ends with **.lab.abc.com**.

If you enter **share3** in the Shares field, mark the Belonging to box, click Storage System and enter **netapp**, the filter adds all the shares that are named **share3** and are exported from the **netapp** storage system.

A selection of filesets that match a pattern

You can add filesets that are associated with file systems, from a storage system to the application or subcomponent. You can select only one resource option per filter creation.

For example, in the Define Name Patterns section, if you enter ***** in the Filesets field, mark the Belonging to box, click Storage System, enter **scale*** and click Save, the filter adds all the filesets that are exposed by storage systems that have names that begin with **scale** to the application.

If you enter ***** in the Filesets field, mark the Belonging to box, click File System and enter **fs1**, the filter adds all the filesets that are associated with the file systems named **fs1** to the application.

A selection of vaults that match a pattern

You can add vaults of a certain name format and vaults that belong to a specific IBM Cloud Object Storage system.

For example, in the Define Name Patterns section, if you enter ***v_cloud*** in the Vault field, select Belonging to, and enter ***.lab.abc.com** in the Cloud Object Storage System field, the filter adds all vaults that have a name that includes **v_cloud** and belong to IBM Cloud Object Storage systems with names that end in **.lab.abc.com**.

If you enter ***v_cloud_ibm** in the Vaults field, select Belonging to, and enter **COS_storage** in the Cloud Object Storage System field, the filter adds all vaults that have a name that ends with **v_cloud_ibm** and belong to the **COS_storage** storage system.


Procedure

To create a filter, complete these steps:

1. In the menu bar, go to Groups>Applications.
2. On the Applications page, right-click a row in the Applications list and click View Details.
3. In the Members section, click Filters.
4. On the Filters page, click Create Filter.
5. On the Create Filter page, select a resource type filter and enter the necessary information.
6. Click Preview to see a list of filter matches before you finalize your filtering criteria.
7. On the Filter Matches page, click Save to create the filter.

Results

The Members section of the details page shows the number of filters that you created, for example, Filters (6). Click Filters to access the Filters page and see the filters that were created. For example, The number of matches: 26, which indicates there are 26 matches for the filter criteria. Click the link to see the resources for the filter matches that are available.

Tip: To view a description of a field on the Create Filter page, hover the mouse pointer over the help icon  for that field.

Previous topic: [Creating applications](#)

Next topic: [Adding subcomponents](#)

Related tasks

- [Creating applications](#)
- [Adding subcomponents](#)

Adding subcomponents

You created a new application and a subcomponent, with their own set of resources and added it to your business hierarchy. You now want to add another level to your business hierarchy.

Procedure

To create a subcomponent for an existing application, complete these steps:

1. In the menu bar, go to Groups>Applications.
2. On the Applications page, select one or more applications, right-click and then, click Add to...>Application.
3. On the Add to Application page, choose the applications that you want to become subcomponents of your selected applications.

- The Hierarchy column shows the path of only the ancestor applications; it does not show departments.
4. Click Save.
 5. Optional: On the details page for the application, click Subcomponents.
 6. On the Subcomponents page, click Create Subcomponent.
 7. On the Create Subcomponent page, enter information including name, description, type, and subtype.
 8. Click Create.

Results

After the subcomponent is created, you can view the results on the Subcomponent page and in the properties. For a more detailed view of the subcomponent, select the subcomponent from the Subcomponent page and click Actions, > View Details.

Previous topic: [Using filters to assign resources to applications](#)

Adding resources manually to applications

You can add resources directly from a resource page to your applications.

About this task

Only the following resources can be added directly to applications:

- Data stores
- Hypervisors
- File Systems (Only server file systems that are monitored through a Storage Resource Agent.)
- Filesets
- Servers
- Shares
- Vaults (Cannot be added with the CLI)
- Virtual Machines (Only can be added with the CLI **modifyappgroup** command)
- Volumes
- Volume Groups

Tip: For any volume that you add to an application, the performance information for that volume, including historical information, is aggregated with the information for the existing volumes.

Note: The application performance feature isn't available to all IBM Spectrum Control users. Look for news from IBM® about when this feature will be available to you.

Procedure

To add a resource directly to an application, complete these steps:

1. On a resource page, right-click a resource and click Add to Application.
For example, to add a server to an application, in the menu bar, go to Groups, > Servers. On the Servers page, select a server from the Servers list, right-click and click Add to Application.
2. On the Add to Application page, select one or more applications to assign the resource to.
3. Click Save.

Results

The resource is added to the application and is displayed on the Applications page.

Related tasks

- [Using filters to assign resources to applications](#)

Adding resources with the command line interface

You can use the **modifyappgroupviafile** command-line interface (CLI) command to automate your bulk assignment of eligible storage resources to applications. To add a single resource to an application, use the **modifyappgroup** command.

Related reference

- [modifyappgroupviafile command](#)
- [modifyappgroup command](#)

Viewing information about applications

You can view information about the application that you created as related resources, filters that are used to populate the application with resources, any subcomponents, capacity trending and performance information.

Procedure


To view the details of an application, use these steps:

1. In the menu bar, go to Groups > Applications.
2. On the Applications page, right-click an application and select View Details.

Results

The details page shows information about the specific application. The information is available from the following sections on the page:

- Overview: The overall capacity trending about the application and performance information. An interactive breadcrumb that shows the hierarchy of applications and departments and a list of all the top level applications.
- Properties: The key information about the application.
- Subcomponents: The number of subcomponents that are a members of the selected application. If there is more than one subcomponent, the number of subcomponents is displayed.
- Members: The information on the resource types that were added directly to the application or subcomponent. The number next to the resource type, for example, Servers (2), shows the number of resources of that type in the application.
- Filters: The number of filters created and the filtering criteria used for the resources that were added to the application.
- Related Resources: The information about the resources that are related to the application.

Tip: To view descriptions of the information that is available on an application details page, click the Help icon  in the upper-right corner of the page.

Related concepts

- [Application and department hierarchies](#)

Related tasks

- [Creating an application hierarchy](#)

Removing applications and subcomponents

You can remove applications and their subcomponents. The data that was collected to monitor the applications is also removed from the database repository.

Procedure

To remove applications and subcomponents from a level of the business hierarchy, complete the following steps:

1. In the menu bar, go to Groups > Applications.
2. On the Applications page, select one or more applications, and click Delete.
 - a. If you select Delete Subcomponents, the subcomponents that belong only to the application are removed.
 - b. If you do not select Delete Subcomponents, any subcomponents of the application, that is being removed, then become members of the parent application or department. If the application is a top level application, and, the subcomponents are not members of any other application, the subcomponents become top level applications.

Results

The applications and subcomponents are removed and are no longer displayed on the Applications page.

To remove subcomponents from within the direct context of the application, use these steps:

1. In the menu bar, go to Groups > Applications.
2. On the Applications page, select an application from the Application list and right-click View Details.
3. Optional: Select a application from the Applications page and right-click View Properties.
 - a. Click the Subcomponents tab and continue with steps 5-6.
4. On the application details page, click Subcomponents in the General section.
5. On the Subcomponents page, select the subcomponent and right-click Remove from Application.
6. Confirm that you want to remove the subcomponent and click OK.

Removing resources from applications

Related resources within the context of an application can be removed from applications. The action is only available when viewing the related resources that belong to an application on the details page.

About this task

The following resources can be removed from an application:

- Volumes
- Shares
- Servers

- File systems (Only server file systems that are monitored through a Storage Resource Agent.)
- Filesets
- Volume Groups
- Hypervisors
- Data Stores

Resources that were added through a filter match are removed if the filter is modified and no longer includes the resource, the filter is removed from IBM Spectrum® Control, or the resource is removed from IBM Spectrum Control. Resources can be added back to the application if the filter is modified and includes the resource that was originally removed, a new filter is created that matches the resource, the resources are added to IBM Spectrum Control, or you use the Add to Application action.

For example, a server named, `abc.xyz.com` is added to an application named, `ApplicationABC` because it matched the server name filter, `abc*`. You decide to remove the `abc.xyz.com` server from the `ApplicationABC` application and remove the `abc.xyz.com` server from IBM Spectrum Control. Later you want to add the `abc.xyz.com` server back to IBM Spectrum Control. The `abc.xyz.com` server is added to the `ApplicationABC` application because the name matches the server resource filter that you originally created for the `ApplicationABC` application. If you do not want the resource added back to the application, you can further refine your filters. Click the View Filters link on the confirmation dialog to make any adjustments to the filtering criteria before creation.

Resources that are automatically added to an application because they are related to resources that were explicitly added to the application *are not* removed from the application. The resources are automatically removed when the corresponding related resource that was added to the application, is also removed from the application.

For example, you add a server to the application that has one volume from a storage system. The volume is added as a related resource of that server in the application. If you attempt to remove the volume from the application, the volume is *not* removed since it is a related resource. The volume is only removed when the server, that was explicitly added to the application, is removed.

Procedure

To remove resources directly from an application, use these steps

1. In the menu bar, go to Groups > Applications.
2. On the Applications page, select an application from the Application list and right-click View Details.
3. On the application details page, click an eligible resource in the Related Resource section.
4. On the resource page, select the resource from the list, and right-click Remove from Application.
5. Confirm that you want to remove the resource.
6. Click Remove.

Results

The resources that were added to the application are removed from the application.

Tip: When you remove a member volume from an application, the performance information for that volume, including historical information, is no longer aggregated with the information for the existing volumes.

Note: The View Performance feature for applications isn't available to all IBM Spectrum Control users. Look for news from IBM® about when this feature will be available to you.

Related tasks

- [Using filters to assign resources to applications](#)

Departments


View information about departments and structure the hierarchy to mirror your business organization. Monitor the detailed information about capacity and space usage that is collected for departments and subdepartments, and monitor the performance of the applications and the application subcomponents that are added to departments and subdepartments.

A department can be hierarchical in its organizational layout. For example, a department might use 15 applications and be part of another five departments. A department might share storage resources with another department, subdepartment or application even if they do not belong in the same hierarchy. For example, in a collaboration scenario, a single IBM® SAN Volume Controller might be shared by multiple departments. You can provide reports on the storage capacity that is used at different levels of your business hierarchy to regulate and compare the capacity storage, and in turn, track the trending of storage usage by specific departments.

The Departments page shows storage capacity information about the top level departments, the subdepartments and any applications that belong to the department that are monitored by IBM Spectrum® Control and the actions that you can use to view and manage the resources that are defined to them. To see more detailed information, view properties and details for the department.

The following information shows the objects that you can add using the Add to Department function.

Table 1. Objects that can be added to the department

Object	Objects that can be added to the department
 Department	<ul style="list-style-type: none"> • Applications • Departments

- [Creating a department hierarchy](#)

To monitor the capacity and space usage of the department, you want to create a department hierarchy and associate applications with the department hierarchy. You can also monitor the performance of the storage resources that are associated with the applications that you add to the department hierarchy.

- [Viewing information about departments](#)

You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments that are used to establish a hierarchical business structure.

- [Removing departments and subdepartments](#)

You can remove departments, subdepartments and applications that are members of the department. The data that is collected about the departments, subdepartments and applications is removed from the database repository.

- [Removing applications from a department](#)

You can remove an application from within a department. This action is only available from the department detail page or the properties for the department.

Creating a department hierarchy

To monitor the capacity and space usage of the department, you want to create a department hierarchy and associate applications with the department hierarchy. You can also monitor the performance of the storage resources that are associated with the applications that you add to the department hierarchy.

About this task

In this scenario, you want to monitor the capacity and space usage of departments in a business organization that sells books. You also want to monitor the performance of the storage resources that are associated with the applications or application subcomponents that are used by the departments.

The sales department of the organization, Book Sales, has these subdepartments:

- Wholesale
- Retail
- Online

To monitor the storage resources that the Books Sales department uses and that each of its subdepartments use, complete these tasks:

- Create a department for Book Sales.
- Create the Wholesale, Retail, and Online subdepartments that you want to add to the department:
 - Create the subdepartments as departments.
 - Add the departments as subdepartments to the department.
 - Add the applications to the subdepartments.

Results

You can now view charts and information about the capacity and space usage of the department and each of the subdepartments on the Departments page and the Subdepartments page. You can also view charts and information about the performance of the storage resources that are associated with the applications or application subcomponents that the departments and subdepartments use.

Tip: In the Subdepartments column for the department on the Departments page, the name of the subdepartment is shown. If you add two or more departments to a subdepartment, the number of subdepartments is shown. To view information about the subdepartments, click the name of the subdepartment or the number.

What to do next

You can view the total capacity and space usage for the department and the capacity and space usage for each of the subdepartments. You can add to your department hierarchy by adding more applications and departments.

1. [Creating departments](#)

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum® Control.

2. [Adding and creating subdepartments](#)

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add an existing subdepartment or create a new subdepartment.

3. [Adding applications to departments](#)

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the Add to Department page become parents of the selected applications.

Creating departments

To model the storage capacity that is consumed in a department for your business environment, in accordance with other department and application members, create departments for monitoring in IBM Spectrum Control.

About this task

Use the web-based GUI to create departments, subdepartments (member departments) and add applications for monitoring different levels of a business hierarchy.

Procedure

To create a department, complete these steps:

1. In the menu bar, go to Groups > Departments.
2. On the Departments page, click Create Department.
3. On the Create Department page, enter information about the department.

You cannot have a duplicate department name on the same level in your department hierarchy. For example, you create a department named *Finance* and create a subdepartment named *Business*. You can create another department named *Business* as a top level department, but if you want to make your new *Business* department a subdepartment of the *Finance* department, you see a message that indicates there is already a department named *Business* that is a member of the *Finance* department; your department creation is denied.

You can customize the icon for your department by clicking the existing icon on the Create Department page and selecting another icon from the Customize Icon page.

4. Click Create and then Close.
5. Optional: On the Add Members page, click Add one or more applications .
 - a. On the populate page, select the applications that to add to your department.
If you do not want to add any of the applications that are listed, click Create Application to create a new application to add to the department.
 - b. Click Finish.
6. Optional: On the Add Members page, click Create subdepartment .
 - a. On the Create Subdepartment page, enter your information.
 - b. Click Create.

Results

To view the newly created department, go to the Departments page.

To see the applications that you added to your department, see the details page and click Applications.

To see the subdepartment you created, see the Subdepartments column on the Departments page or see the details page and click Subdepartments.

What to do next

Continue to add to your department hierarchy by adding applications and departments.

Next topic: [Adding and creating subdepartments](#)

Related tasks

- [Adding and creating subdepartments](#)
- [Adding applications to departments](#)

Adding and creating subdepartments

To further enhance your data modeling at the department level and to add to your overall business hierarchy for monitoring, add an existing subdepartment or create a new subdepartment.

Procedure

To add a subdepartment to the department level hierarchy, use these steps:

1. In the menu bar, go to Groups > Departments.
2. On the Departments page, select one or more departments, right-click and click Add to Department.
3. On the Add to Department page, select one or more departments from the list and click Save.
The department or departments you select from the Add to Department page become parents of the selected department. The Hierarchy column shows the longest ancestry path of the department. You can also view the applications that are members of the department on the Departments page.
4. Optional: Create a new subdepartment:
 - a. Click Create Department to open the Create Department page.
 - b. Enter information about the department.
 - c. Click Create.
 - d. On the Add Members page, click Create a subdepartment.
 - e. On the Create Subdepartment page, enter your information.
 - f. Click Create.
5. Optional: Create a new subdepartment from the details page of the existing department:
 - a. In the General section, click Subdepartments to open the Subdepartments page.
 - a. Click Create Subdepartment to open the Create Subdepartment page.
 - b. Enter information about the subdepartment.
 - c. Click Create.

Results

To see the subdepartment you created, see the Subdepartments column on the Departments page or click Subdepartments on the details page.

What to do next

You can continue to create or add subdepartments to the department or add applications to the department.

Previous topic: [Creating departments](#)

Next topic: [Adding applications to departments](#)

Related tasks

- [Creating departments](#)

- [Adding applications to departments](#)

Adding applications to departments

To further enhance the data storage modeling in a business environment and see the storage capacity that is used at the different levels of your business hierarchy, add applications to departments. The departments you select from the Add to Department page become parents of the selected applications.

Procedure

To add applications to departments, complete these steps:

1. In the menu bar, go to Groups > Applications.
2. On the Applications page, select one or more applications, right-click and then, click Add to... > Department.
3. On the Add to Department page, select one or more departments to add to the selected application or applications.
The Hierarchy column shows the longest ancestry path of the department.
4. Click Save.
5. Optional: On the Create Department page, enter information about the department.
6. Click Create.
7. On the Add Members page, click Add one or more applications.
8. On the populate page, select the applications that to add to your department.
If you do not want to add any of the applications that are listed, click Create Application to create a new application to add to the department.
9. Click Finish.

Results

View the added department or departments in the Departments column on the Applications page or the Departments tab of the application properties.

Previous topic: [Adding and creating subdepartments](#)

Related tasks

- [Creating applications](#)
- [Creating departments](#)

Viewing information about departments

You can view detailed information about departments for use in storage capacity monitoring and performance troubleshooting. You can view additional information on the applications and subdepartments that are used to establish a hierarchal business structure.

Procedure


To view the details of a department, use these steps:

1. In the menu bar, go to Groups > Departments.
2. On the Departments page, right-click a department in the list and select View Details.

Results

The department details page shows information about the specific department. The information is available in the following sections on the page:

- Overview: The overall capacity trending and performance information about the department. An interactive breadcrumb that shows the hierarchy of the departments and applications and a list of all the top level departments.
- Properties: The key information about the department.
- Applications: The applications that belong to the department. If there is more than one application, the number of applications is displayed.
- Subdepartments: The subdepartments that belong to the department. If there is more than one subdepartment, the number of subdepartments is displayed.
- Related Resources: Information about related resources in the storage environment that the applications, which belong to the department, are using or are connected to.

Tip: To view descriptions of the information that is available on a department details page, click the Help icon  in the upper-right corner of the page.

Related concepts

- [Application and department hierarchies](#)

Related tasks

- [Creating a department hierarchy](#)

Removing departments and subdepartments

You can remove departments, subdepartments and applications that are members of the department. The data that is collected about the departments, subdepartments and applications is removed from the database repository.

Procedure

To remove departments, subdepartments, and applications, that are members of the department, from a level of the business hierarchy, complete these steps:

1. In the menu bar, go to Groups > Departments.
2. On the Departments page, select one or more departments, right-click and click Delete.
3. Confirm the information and click Delete.
 - a. If you select Delete subdepartments and applications, the subdepartments and applications that belong only to the department, that is being removed, are removed.
 - b. If you do not select Delete subdepartments and applications, the subdepartments and applications of the department that are being removed become members of the parent department. If the department is a top level department, and the subdepartments are not members of any other department, they become top level departments.

Results

The departments, subdepartments and applications are removed from the business level hierarchy and are no longer displayed on the Departments page. To remove subdepartments from within the direct context of a department, use these steps:

1. In the menu bar, go to Groups > Departments.
2. On the Departments page, select a department from the list and right-click View Details.
3. Optional: Select a department from the Departments page and right-click View Properties.
 - a. Click the Subdepartments tab and continue with steps 5 and 7.
4. On the department details page, click Subdepartments in the General section.
5. On the Subdepartments page, select the subdepartment and right-click Remove from Department.
6. Optional: On the Subdepartments page, select the subdepartment and click Delete.
7. Reply to the confirmation message and click OK.

Removing applications from a department

You can remove an application from within a department. This action is only available from the department detail page or the properties for the department.

Procedure

To remove applications from within the context of a department, use these steps:

1. In the menu bar, go to Groups > Departments.
2. On the Departments page, select a department from the list and right-click View Details.
3. Optional: Select a department from the Departments page and right-click View Properties.
 - a. Click the Applications tab and continue with steps 5-7.
4. On the department details page, click Applications from the General section.
5. On the Applications page, select the application and right-click Remove from Department.
6. Confirm the information.
7. Click OK.

Results

The applications that were added within the context of a department are eligible for removal from within department.

Related tasks

- [Viewing information about departments](#)

Application and department hierarchies

The interactive breadcrumb on the Overview page provides a visual representation of the hierarchical relationships for applications and departments and enables you to navigate quickly between the elements in the hierarchy.

The breadcrumb uses default or customized icons to represent the related applications and departments. The greater-than sign (>) serves as the hierarchy separator and ends with the name of the selected application or department. Only the ancestor names of the selected application are shown.



Overview

Multiple (2) V DB2 > Data
Epic
Records

When you move your mouse pointer over an icon in the breadcrumb, you see the siblings for that application or department. If you move your mouse pointer, for example, over the Data subcomponent, the related application and department elements are shown.

The visual representation of the hierarchy stops when either there are no additional parents or multiple parents are encountered from the selected application or department. When multiple parents are encountered, the number of parents are shown in parenthesis. For example, *Multiple (2)*. Moving the mouse pointer over *Multiple* shows a list of the parents. The siblings of multiple parents are *not* shown.

Move your mouse pointer over a sibling in the breadcrumb and it becomes a name link. Click the application or department name link to navigate to the details for that application or department and view the breadcrumb if it exists.

You can add resources directly on the Overview page for applications from the breadcrumb. For example, in this case, you create an application named Db2® and create two subcomponents, Logs and Data. On the Overview page for the Logs subcomponent, click the Logs subcomponent name link in the breadcrumb and click Filters to create a resource filter for that subcomponent. Then, click the Data subcomponent name link and create a resource filter for that application subcomponent. To see a visual representation of the hierarchy for applications, subcomponents, departments, and subdepartments click View Details.

Related tasks

- [Creating an application hierarchy](#)
- [Creating a department hierarchy](#)

Exporting information about applications and departments

You can save information about applications and departments to a PDF, CSV, or HTML file. Information that you can export from the GUI includes all the values that are being shown in the columns for a list of applications and departments.

Before you begin

The information that you export to a file is organized according to the sorting, filtering, and column order that is defined for a list. Before you export information, complete the following tasks to configure a list:

- Change the order of columns in the list as you want them to appear in the generated file.
- Hide columns in the list that you do not want to include in the generated file. Information in hidden columns is not exported.
- Sort the rows in the list as you want them to appear in the generated file. Each list in the GUI has a column or set of columns that determines the order of its rows.
- Filter the list to show only the applications and departments that you want to export. Use filtering to limit the rows that are shown in a list based on value in a specific column or set of default columns.

About this task

You can export information that is shown on the application and department list and detail pages.

Applications and departments list and detail pages

You can export information about the top level applications and departments and the related resources that are shown on application and department list and detail pages.

For example, you can export information about monitored applications and their subcomponents and departments on the Applications page. When you select the export action, the generated file includes the column values for each monitored application in the list.

For example, you can export information about the servers that are associated with an application on the details page.

Procedure

1. In the GUI, go to the Applications or Departments page and select applications or departments for which you want to export information.
2. Click Actions and select Export. *file_format*, where **file_format** represents the format that you want to use for a file.

The following formats are available:

CSV (comma-separated values)

A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information in a CSV file has the following format:

```
"Column_name", "Column_name", "Column_name", "Column_name", "Column_name", "Column_name"
"data", "data", "data", "data", "data", "data"
```

where **Column_name** represents the name of a column in a table and **data** represents the data that is associated with a column.

PDF

You can view a PDF file with Acrobat reader. The information that you export to a PDF file is formatted into a table.

Tip: When you export a table of data that contains many columns, the rows in that table might span multiple pages in the resulting PDF. For example, if a table contains 20 columns, then the row for a specific resource might be shown on more than one page (10 columns on the first page and 10 columns on the second page).

To reduce the number of pages that a table spans, before you select the export action, hide the columns that you do not want to include in the PDF. Continue hiding columns until the table no longer spans multiple pages.

HTML

You can view an HTML file with a web browser. The information that you export to an HTML file is formatted into a table.

3. Optional: Depending on how your web browser is configured, you can specify the name and location of the generated file.
4. Click the save option to export the information to a file.

Related concepts

- [Customizing lists](#)

General groups

View and administer the general groups that are used to alert on logically related storage resources. You can define alerts for a general group to notify you about changes in the configuration, attributes, and performance of the resources in the group.

Organizing your resources into general groups can be helpful in the following situations:

- When you want to receive alert notifications about changes for a subset of the resources of a particular type. For example, you want to detect when the ports that are used for replication on your SAN Volume Controller have insufficient buffer-to-buffer credit. You do not want the alerts to apply to other ports on your SAN Volume Controller. You can group the ports that are used for replication and then define alerts for the group.
- When you want to receive alert notifications about changes for a group of resources that are logically related. Examples include all the servers that use a particular operating system or all the storage systems at a specific location.

General group hierarchies

Organizing resources into general groups and their subgroups can be helpful when you want to quickly view information about a group of resources, but you also want to view information about subgroups of resources within the group.

Only top-level groups in the general group hierarchy are shown on the General Groups page. If the general group has a subgroup, the name of the subgroup is shown in the Subgroups column. If the general group has more than one subgroup, click the number of groups in the Subgroups column to view the subgroups.

- **[Creating a general group hierarchy](#)**
Organizing resources into general groups and their subgroups can be helpful when you want to quickly view information about a group of resources, but you also want to view information about subgroups of resources within the group.
- **[Adding resources to general groups](#)**
Add resources to general groups so that you can receive alert notifications about changes in the configuration, attributes, and performance of the resources in the group. You can add resources to one or more existing groups or you can create a new group.
- **[Viewing and administering general groups](#)**
View and administer the general groups that you use to alert on logically related storage resources.

Related tasks

- [Defining alerts for general groups](#)

Creating a general group hierarchy

Organizing resources into general groups and their subgroups can be helpful when you want to quickly view information about a group of resources, but you also want to view information about subgroups of resources within the group.

About this task

In this scenario, you want to monitor all the ports on your SAN Volume Controller but you also want to separately monitor the following subsets of ports on the SAN Volume Controller:

- Ports that are used for inter-node communication
- Ports that are used for host I/O exchanges

To add the ports to a general group hierarchy, complete the following steps:

1. Add the ports that are used for inter-node communication to a general group, SVC Inter Node Ports.
2. Add the ports that are used for host I/O exchanges to another general group, SVC Host I/O Ports.
3. Add the groups as subgroups of a parent general group, SVC All Ports.

You can now quickly view information about the resources in the hierarchy on the details page for the general groups.

When you view the SVC Inter Node Ports or SVC Host I/O Ports groups, you see information about the specific ports that are members of that subgroup.

When you view the parent group, SVC All Ports, you see information about the ports that are members of the parent group and also the ports in the subgroups.

- **[Creating general groups and adding resources](#)**
When you create and associate storage resources with a general group, you can quickly view information about the group and its resources.
- **[Adding general groups as subgroups](#)**
When you add one or more general groups as subgroups to another general group, you can quickly view information about the resources in the group hierarchy.

Creating general groups and adding resources

When you create and associate storage resources with a general group, you can quickly view information about the group and its resources.

About this task

You can add the following resources and their internal resources to a general group:

- Storage systems
- Servers
- Hosts
- Hypervisors
- Switches
- Fabrics
- Other general groups

Only resources that you specifically add to a group are included as members of the group. For example, if you add a SAN Volume Controller, the internal resources of the SAN Volume Controller, such as volumes and pools, are not automatically added to the group. To add the volumes and pools, you must specifically select those resources and add them to the group.

Procedure

To create a general group and add resources to the group, complete these steps:

1. Take one of the following actions to go to the list page for the resource or group that you want to add:
 - To add top-level resources such as storage systems, servers, hypervisors, switches, or fabrics, go to the appropriate resource page. For example, to add servers, in the menu bar, click Servers >> Servers.
 - To add internal resources of top-level resources, complete the following steps:
 - Go to the resource list page for the top-level resource. For example, to add internal resources of a SAN Volume Controller, in the menu bar, click Storage >> Block Storage Systems.
 - Right-click the resource and click View Details.
 - In the Internal Resources section, click the type of resource that you want to add, for example, volumes or pools.
 - To add general groups, in the menu bar, click Groups >> General Groups.
2. Right-click one or more resources or groups and click Add to General Group.
3. Click Add to new group and specify a name and description for the new group.
4. Optional: Customize the icon for a new group by clicking the existing icon and selecting another icon.
5. Click Save.

Results

You can now quickly view information about the resources in the general group hierarchy on the details page for the general group.

Related tasks

- [Adding resources to general groups](#)

Related reference

- [Viewing and administering general groups](#)

Adding general groups as subgroups

When you add one or more general groups as subgroups to another general group, you can quickly view information about the resources in the group hierarchy.

Procedure

1. In the menu bar, click Groups >> General Groups.
2. Right-click one or more general groups and click Add to General Group.
3. Take one of the following actions:
 - To add the groups as subgroups of existing groups, click Add to existing groups, click one or more groups in the list, and click Save.
 - To add the groups as subgroups of a new group, click Add to new group, specify a name and description for the new group, and click Save.

Related tasks

- [Adding resources to general groups](#)

Related reference

- [Viewing and administering general groups](#)

Adding resources to general groups

Add resources to general groups so that you can receive alert notifications about changes in the configuration, attributes, and performance of the resources in the group. You can add resources to one or more existing groups or you can create a new group.

About this task

The following resources and their internal resources can be added to a general group:

- Storage systems*
- Hosts
- Hypervisors
- Switches
- Fabrics
- Other general groups

Tip: * For IBM® Cloud Object Storage, only vaults can be added to general groups.
Only resources that you specifically add to a group are included as members of the group. For example, if you add a SAN Volume Controller, the internal resources of the SAN Volume Controller, such as volumes and pools, are not automatically added to the group. To add the volumes and pools, you must specifically select those resources and add them to the group.

You can create a general group hierarchy by adding one or more general groups as subgroups of a general group. When you define a general group alert for a resource such as a volume, the alert applies to all the volumes that belong to that group and all its subgroups.

Procedure

To add resources to a general group, complete these steps:

1. Take one of the following actions to go to the list page for the resource or group that you want to add:
 - To add top-level resources such as storage systems, servers, hypervisors, switches, or fabrics, go to the appropriate resource page. For example, to add servers, in the menu bar, click Servers. >> Servers.
 - To add internal resources of top-level resources, complete the following steps:
 - Go to the resource list page for the top-level resource. For example, to add internal resources of a SAN Volume Controller, in the menu bar, click Storage. >> Block Storage Systems.
 - Right-click the resource and click View Details.
 - In the Internal Resources section, click the type of resource that you want to add, for example, volumes or pools.
 - To add general groups, in the menu bar, click Groups. >> General Groups.
2. Right-click one or more resources or groups and click Add to General Group.
3. Take one of the following actions:
 - To add the resources or groups to a new group, click Add to new group and specify a name and description for the new group.
 - To add the resources or groups to one or more existing groups, click Add to existing groups and click the appropriate groups in the list.
4. Optional: Customize the icon for a new group by clicking the existing icon and selecting another icon.
5. Click Save.

What to do next

To view the resources and subgroups that are members of the group, complete the following steps:

1. Go to Groups. >> General Groups.
2. Right-click the general group and select View Details.
On the general group details page, all the different resource types, such as volumes or pools, that are members of the group or members of its subgroups are shown in the Members section.
3. Click the resource type to view the specific resources. For example, click Volumes to view all the volumes that are in the general group or in its subgroups.
4. Click Subgroups to view the child groups.

To define alerts for the general group, use the View Alert Definitions action on the General Groups page.

Related tasks

- [Creating general groups and adding resources](#)
- [Adding general groups as subgroups](#)

Related reference

- [Viewing and administering general groups](#)

Viewing and administering general groups

View and administer the general groups that you use to alert on logically related storage resources.

Administering general groups

Table 1. Administering general groups

Actions	Navigation
Create a general group and add resources.	1. In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add file storage systems, click Storage. >> File Storage Systems. 2. Right-click one or more resources and click Add to General Group. 3. Click Add to new group, specify a name and description for the new group, and click Save.

Actions	Navigation
Add resources to general groups.	<ol style="list-style-type: none"> 1. In the menu bar, click the resource list page for the resources that you want to add. For example, if you want to add block storage systems, click Storage_>Block Storage Systems. 2. Right-click one or more resources and click Add to General Group. 3. Take one of the following actions: <ol style="list-style-type: none"> a. To add the resources to one or more existing groups, click Add to existing groups, click the appropriate groups in the list, and click Save. b. To add the resources to a new group, click Add to new group, specify a name and description for the new group, and click Save.
Set which alert policy manages a general group Learn more	<ol style="list-style-type: none"> 1. Go to the details page for the general group for which you want to set the policy. In the menu bar, click Groups_>General Groups. 2. Right-click one of the general groups, then click View Details. 3. Click Alerts in the General section. 4. Click Set Policy from the Policy Actions menu. <p>To set the alert policy for multiple general groups, go to Settings_>Alert Policies. Double-click the policy, click the resources tab, then click Edit Resources.</p>
Create an alert policy from the alert definitions and notification settings in a general group	<ol style="list-style-type: none"> 1. Go to the details page for the general group from which you want to create the policy. In the menu bar, click Groups_>General Groups. 2. Right-click one of the general groups, then click View Details. 3. Click Alerts Definitions in the General section. 4. Click Create Policy from the Policy Actions menu.
View or modify the alert policy that manages a general group Learn more about defining alerts and notification settings	<ol style="list-style-type: none"> 1. Go to the details page for the general group whose policy you want to view. In the menu bar, click Groups_>General Groups. 2. Right-click one of the general groups, then click View Details. 3. Click Alerts in the General section. 4. Click View Policy from the Policy Actions menu.
Delete general groups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups_>General Groups. 2. Right-click one or more general groups and click Delete. 3. To delete subgroups, ensure that the Delete subgroups? check box is selected. <p>When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product.</p>
Remove resources from a general group.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups_>General Groups. 2. Right-click a general group and click View Details. 3. In the Members section, click the type of resource that you want to remove. For example, if you want to remove some of the volumes from the group, click Volumes. 4. Right-click the resources that you want to remove and click Remove from General Group. 5. Click OK to confirm that you want to remove the resources from the general group and from all of its subgroups.
Add general groups as subgroups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups_>General Groups. 2. Right-click one or more general groups and click Add to General Group. 3. Take one of the following actions: <ol style="list-style-type: none"> a. To add the groups as subgroups of existing groups, click Add to existing groups, click one or more groups in the list, and click Save. b. To add the groups as subgroups of a new group, click Add to new group, specify a name and description for the new group, and click Save.
Remove subgroups from a group hierarchy.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups_>General Groups. 2. Right-click a general group and click View Details. 3. Click Subgroups in the General section. 4. Right-click the general groups that you want to remove and click Remove from General Group. <p>When you remove a subgroup from its parent group, the subgroup is moved to the same level in the hierarchy as the parent group. The subgroup is still monitored by the product.</p>
Create a subgroup.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups_>General Groups. 2. Right-click the general group that you want to create a subgroup for and click View Details. 3. Click Subgroups in the General section. 4. Click Create Subgroup. 5. Specify a name and description for the new group. 6. Customize the icon for the group by clicking the existing icon and selecting another icon. 7. Click Create.
Delete subgroups.	<ol style="list-style-type: none"> 1. In the menu bar, click Groups_>General Groups. 2. Right-click the general group that contains the subgroups that you want to delete and click View Details. 3. Click Subgroups in the General section. 4. Right-click one or more subgroups and click Delete. 5. To delete further subgroups, ensure that the Delete subgroups? check box is selected. <p>When a group is deleted, the information that is associated with the group, such as alerts and alert definitions, is removed from the product. The resources that were members of the group are removed from the group but are still monitored by the product.</p>

Table 2. Viewing general groups

Actions	Navigation
View a list of general groups.	<p>In the menu bar, click Groups > General Groups.</p> <p>The following general groups are shown:</p> <ul style="list-style-type: none"> The top-level groups in general group hierarchies All general groups that are not part of a general group hierarchy
View subgroups in a general group hierarchy.	<ol style="list-style-type: none"> In the menu bar, click Groups > General Groups. Click the number of groups or the group name that is shown in the subgroups column. Click the Subgroups tab.
View details about a general group, including the alert notifications, the resources that are members of the group, and the subgroups.	<ol style="list-style-type: none"> In the menu bar, click Groups > General Groups. Right-click a general group and select View Details.
View the alert definitions and notification settings for a general group.	<ol style="list-style-type: none"> In the menu bar, click Groups > General Groups. Right-click a general group and select View Alert Definitions.
View the alerts that were triggered for a general group.	<ol style="list-style-type: none"> In the menu bar, click Groups > General Groups. Right-click a general group and select View Details. Click Alerts in the General section.
View reports about general groups.	<p>In the menu bar, click Cognos. The Cognos® Analytics reporting tool is included in Tivoli® Common Reporting.</p> <p>To view reports about general groups, view reports about storage resource groups in the Cognos Analytics reporting tool.</p>

Related tasks

- [Creating general groups and adding resources](#)
- [Adding resources to general groups](#)
- [Adding general groups as subgroups](#)

Provisioning storage

You can assign storage to servers, hypervisors, and clusters on servers and hypervisors in the GUI.

About this task

Discontinued support: Provisioning is no longer supported in IBM Spectrum® Control. While the feature might still work in this release, it's recommended that you use another tool for your provisioning needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

You can also assign the following types of storage to servers, hypervisors, and clusters on servers and hypervisors:

- Volumes
- NAS file shares

When you assign storage, you must associate the storage with a service class. You can associate the storage with a default service class such as gold, silver, or bronze, or you can create a user-defined service class. You can also create capacity pools to restrict the provisioning of storage to a defined set of storage resources such as pools on a storage system in a particular location.

- [Configuring IBM Spectrum Control for provisioning](#)
- [Provisioning storage with the IBM Spectrum Control GUI](#)

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum Control.

Configuring IBM Spectrum Control for provisioning

Before you can use the simplified provisioning capabilities that are provided with IBM Spectrum® Control, you must configure your environment.

Before you begin

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

About this task

To provision storage, you are required to specify only the storage capacity and storage quality that is required. After volumes are created, IBM Spectrum Control can analyze and optimize volume performance.

Procedure

To configure IBM Spectrum Control for provisioning, complete the following tasks:

1. Assign storage pools to tiers by setting the tier level of each storage pool. The tier levels are number tags that can reflect any tier structure in the environment. You must assign storage pools to tiers to later provision volumes that require a certain tier level and to optimize the placement of volumes. Storage pools do not need to be assigned to tiers to provision volumes that do not require a tier level, or to provision file shares.
 2. Create service classes. A service class is a set of properties that describe capabilities and characteristics of storage resources. A service class typically describes a particular quality of service, and is used during provisioning to describe storage requirements. For example, a block-storage service class specifies properties such as a required RAID level, and whether storage resources must be able to encrypt or thin provision volumes.
 3. Optional: Create capacity pools. You can restrict provisioning requests to a capacity pool. Capacity pools are groups of storage resources. You can use capacity pools to separate storage resources in any way that serves the needs of your environment or business. For example, a capacity pool might contain the storage resources that are allocated to a particular department or division of your business, or to a particular application.
- [Configuring service classes](#)
Configure service classes to represent the different levels or types of storage quality that you want to provision. When you are configuring service classes for your installation, consider whether certain storage requests must always be satisfied from a particular set of resources. Also, consider whether you want to allow users without administrator privileges to provision storage, and whether you want their provisioning requests to require administrator approval.
 - [Configuring capacity pools](#)
Configure capacity pools to track the used and available capacity for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.
 - [Block storage: Calculating available capacity and determining the placement of volumes](#)
For each block-storage service class, IBM Spectrum Control calculates the amount of available capacity for volumes of that service class. When you request volumes of a particular service class, IBM Spectrum Control recommends the best location for the volumes from the available capacity.
 - [File storage: Calculating available capacity and determining the placement of shares](#)
For each file-storage service class, IBM Spectrum Control calculates the amount of available capacity for NAS file shares of that service class. When you request file shares of a particular service class, IBM Spectrum Control recommends the best location for the shares from the available capacity.
 - [Changing the default host definition for provisioned storage](#)
Use the `setdscfg` command to change the default host definition for fibre-channel ports when provisioning storage from resources that run IBM Spectrum Virtualize.

Configuring service classes

Configure service classes to represent the different levels or types of storage quality that you want to provision. When you are configuring service classes for your installation, consider whether certain storage requests must always be satisfied from a particular set of resources. Also, consider whether you want to allow users without administrator privileges to provision storage, and whether you want their provisioning requests to require administrator approval.

Before you begin

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum® Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

- [Creating service classes](#)
Create service classes to specify storage requirements for future provisioning requests. A service class describes attributes and capabilities of block storage or file storage resources, and typically describes particular quality of service.
- [Viewing the details of service classes](#)
You can view the provisioning requirements and storage constraints that a service class defines. You can also view the volumes or shares that were provisioned by using the service class.
- [Modifying service classes](#)
You can modify the attributes of a service class to change the provisioning requirements, capacity pool storage constraints, user permissions, and service class name.
- [Deleting service classes](#)
You can delete a service class that is no longer needed.
- [Tagging resources to satisfy custom requirements](#)
A service class can specify tags to create custom requirements for provisioning. Add the same tags to storage resources that satisfy the custom requirements. When the service class is specified during provisioning, only the tagged resources are candidates for provisioning.

Creating service classes

Create service classes to specify storage requirements for future provisioning requests. A service class describes attributes and capabilities of block storage or file storage resources, and typically describes particular quality of service.

Before you begin

You must have Administrator privileges to create a service class.

About this task

The IBM Spectrum® Control GUI guides you through the steps to create a service class. When you provision a volume or share, you define storage requirements by specifying a service class. A block-storage service class specifies the requirements for the block-storage resource from which a new volume is allocated. A file-storage service class defines requirements for the file storage resources from which a new share is allocated.

You can associate a service class with one or more capacity pools. If you do, future provisioning requests can be satisfied only by storage resources in the specified capacity pools.

You can selectively grant users who do not have administrator privileges permission to provision storage by using the service class.

Restrictions:

- To create a block-storage service class, you must have the IBM Spectrum Control Advanced Edition license.
- To associate the service class with a capacity pool, one or more capacity pools must be defined in the IBM Spectrum Control database. If no capacity pool is defined, the Specify Capacity Pool page is not displayed.

Procedure

Start this task at the Service Classes page. To open the Service Classes page, select Advanced Analytics > Cloud Configuration > Work With Service Classes. To create a service class, complete the following steps:

1. From the Service Classes page, click Create Service Class.
2. Select the type of service class that you want to create:
 - To create a block-storage service class, click the Block icon.
 - To create a file-storage service class, click the File icon.
3. In the Define Properties page, complete the following steps:
 - a. Specify a name for the service class.
 - b. Optional: Specify a service class description.
Descriptions are displayed in the table of service classes on the Service Classes page.
 - c. Specify the general attributes of the service class.
Most of the general properties specify requirements for provisioning. Some of the general properties specify configuration instructions for the storage resources or for the volumes or shares that are provisioned from them.
Tip: To display help information about a service class property, move the mouse pointer over the field or control. Then, move the mouse pointer over the question mark icon displayed next to the field.
4. Optional: Specify or modify the advanced properties of the service class.
 - a. From the Define Properties page, click Advanced.
 - b. If you are creating a block-storage service class, and thin provisioning is enabled for the service class, specify configuration instructions for thin provisioning.
 - c. If you are creating a block-storage service class, select how the multipath driver on the host is configured. For hypervisors, the multipath driver is configured only if provisioning is initiated from the vSphere Web Client extension for IBM Spectrum Control.
 - d. Create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.
5. Optional: In the Specify Capacity Pool page, associate the service class with one or more capacity pools.
If capacity pools are associated with the service class, future provisioning requests that specify the service class can be satisfied only by storage resources in the capacity pools.
6. Optional: Grant permission to individual users, who do not have administrator privileges, to provision storage by using the service class.
By default, users must be assigned to the Administrator role to provision storage. However, you can selectively grant users who are limited to the Monitor or External Application role permission to provision by using this service class. You can also specify whether administrator approval is required to run provisioning tasks that are created by these users.
7. Click Finish
The new service class is listed in the table of service classes on the Service Classes page, and can be specified in future provisioning requests.

Viewing the details of service classes

You can view the provisioning requirements and storage constraints that a service class defines. You can also view the volumes or shares that were provisioned by using the service class.

About this task

From the Service Classes page, you can open a properties page to view the following information for a service class:

- The properties of a service class, such as the required RAID level or storage tier for a block-storage service class.
- The storage constraints that are imposed by the service class. The service class can allow provisioning from all available storage, or can limit provisioning to one or more capacity pools.
- The users who can provision storage by using the service class.
- The volumes or shares that were provisioned by using the service class.
- The candidate storage for the service class. The *candidate storage* is the set of storage resources that can provide the service class.

Procedure

To open the service class properties window from the Service Classes page, complete the following steps:

1. In the menu bar in the web-based GUI, go to Advanced Analytics > Cloud Configuration, and click Work With Service Classes.
The Service Classes page is displayed.
2. From the list of service classes, select the service class, and then select Actions > View / Modify.
The service class properties window is displayed.

Modifying service classes

You can modify the attributes of a service class to change the provisioning requirements, capacity pool storage constraints, user permissions, and service class name.

Before you begin

You must have Administrator privileges to modify a service class.

About this task

Important: When volumes or shares are provisioned by specifying a service class, they are associated with the service class. The Volumes page and the Shares page in the web-based GUI can display the name of the service class that was used to provision the volume or share. If you modify a service class, the volumes or shares are still associated with the service class, but might be on storage resources that no longer satisfy the requirements of the service class. Depending on your changes to the service class, users might incorrectly assume that the new volumes or shares have attributes that they do not possess.

Procedure

Start this task at the Service Classes page. To open the Service Classes page, select **Advanced Analytics** > **Cloud Configuration** > **Work With Service Classes**. To modify a service class, complete the following steps:

1. From the list of service classes, select a service class and then select **Actions** > **View / Modify**.
The service class properties window is displayed.
2. Optional: Click the General tab to modify the general attributes of the service class.
Most of the general properties specify requirements for provisioning. Some of the general properties specify configuration instructions for the storage resources or for the volumes or shares that are provisioned from them.
Tip: To display help information about a service class property, move the mouse pointer over the field or control. Then, move the mouse pointer over the question mark icon displayed next to the field.
3. Optional: Specify or modify the advanced properties of the service class.
 - a. From the General tab, click **Advanced**.
 - b. If you are modifying a block-storage service class, and thin provisioning is enabled for the service class, specify configuration instructions for thin provisioning.
 - c. If you are modifying a block-storage service class, select how the multipath driver on the host is configured.
 - d. Create custom requirements for the service class by specifying up to three custom tags. To provide the service class, storage resources must have all the same tags that are specified in the service class.
4. Optional: Modify storage constraints.
If one or more capacity pools are defined in the IBM Spectrum® Control database, you can associate the service class with one or more capacity pools. If capacity pools are associated with the service class, future provisioning requests that specify the service class can be satisfied only by storage resources in the capacity pools. If a capacity pool was previously associated with the service class, you can disassociate it from the service class.
5. Optional: Modify user permissions.
By default, users must be assigned to the Administrator role to provision storage. However, you can selectively grant users who are limited to the Monitor or External Application role permission to provision by using this service class. You can also specify whether administrator approval is required to run provisioning tasks that are created by these users.
If users in the Monitor or External Application roles have permission to provision by using the service class, you can revoke that permission. You can specify that only administrators can provision by using the service class, or you can selectively remove users. To remove a user, right-click the user name in the table and then select **Remove** from the context-sensitive menu.

Deleting service classes

You can delete a service class that is no longer needed.

Before you begin

You must have Administrator privileges to delete a service class.

About this task

Deleting a service class does not affect the storage volumes or NAS shares that are already provisioned by using the service class. However, the storage volumes or NAS shares are no longer associated with any service class. When a volume that is not associated with a service class is optimized, the volume can be migrated to any available storage pool. The storage pool to which the volume is migrated might not satisfy the requirements of the service class that was specified when the volume was provisioned.

Procedure

To delete a service class from the Service Classes page, complete these steps:

1. In the menu bar, go to **Advanced Analytics** > **Cloud Configuration**, and click **Work With Service Classes**.
The Service Classes page is displayed.
2. From the list of service classes, select the service class, and then select **Actions** > **Delete**.

Tagging resources to satisfy custom requirements

A service class can specify tags to create custom requirements for provisioning. Add the same tags to storage resources that satisfy the custom requirements. When the service class is specified during provisioning, only the tagged resources are candidates for provisioning.

Before you begin

You must have Administrator privileges to tag storage resources.

About this task

When you provision storage, you indicate your storage requirements by specifying a service class. Certain properties of the service class describe capabilities that storage resources must have to be a candidate for provisioning. In addition to the default properties, a service class can define custom requirements by specifying up to three custom tags. To provide the service class, a storage resource must have all the same tags as the service class. The ability to define custom requirements by tagging service classes and storage resources gives you more control over the selection of storage resources during provisioning. By using tags, you can filter the candidates for provisioning by using any criteria that you choose.

When you provision volumes, you specify the requirements by using a block-storage service class. Based on the requirements of the service class, IBM Spectrum® Control identifies a storage pool for the volume. If the service class specifies tags, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

When you provision shares, you specify the requirements by using a file-storage service class. Based on the requirements of the service class, IBM Spectrum Control identifies a file system or Network Shared Disk (NSD) for the share. If the service class specifies tags, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

Procedure

Tag only the storage resources that satisfy the intended custom requirements of the service class. During provisioning, only the tagged resources can provide the service class.

1. Open the resource list page for the type of storage resource that you want to tag. To satisfy the custom tag requirements of a block-storage service class, you can tag a storage pool or its containing block storage system. To satisfy the custom tag requirements of a file-storage service class, you can tag a file system or NSD, or the containing file storage system.

Important: The tags on a containing storage system apply to an internal pool, file system, or NSD resource only if the internal resource is not tagged.

- If you are tagging a storage pool, from the menu bar, go to Storage > Pools.
 - If you are tagging the block storage system that contains the pools, complete the following steps:
 - In the menu bar, go to Storage > Block Storage Systems.
 - Click the Block Storage tab.
 - If you are tagging a file system or NSD, complete the following steps:
 - In the menu bar, go to Storage > Block Storage Systems.
 - Click the File Storage tab.
 - From the list of storage systems, select the storage system that contains the file system or NSD, and then select View Details. Alternatively, you can select Actions > View Details.
 - In the Internal Resources area, click File Systems or Network Shared Disks.
 - If you are tagging the file storage system that contains the file systems or NSDs, complete the following steps:
 - In the menu bar, go to Storage > Block Storage Systems.
 - Click the File Storage tab.
2. In the table of storage systems, pools, file systems, or NSDs, select the storage resource that you want to tag, and then select View Properties. Alternatively, you can select Actions > View Properties.
 3. In the properties notebook for the storage system, pool, file system, or NSD, click Edit.
 4. In the custom tag fields of the properties notebook, specify tags that match the tags that are specified in the service class. To provide the service class, the resource must have all of the same tags.
 5. In the properties notebook, click Save.

Configuring capacity pools

Configure capacity pools to track the used and available capacity for block and file storage on any set of storage resources. Provisioning requests can also be restricted to resources in a capacity pool.

Before you begin

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum® Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

- **[Adding resources to capacity pools](#)**
Divide storage resources into capacity pools to later restrict provisioning or optimization to a specific set of resources. You can also use capacity pools to track the used and available capacity for block and file storage on any set of storage resources. You can add resources to an existing capacity pool, or you can create a new capacity pool.
- **[Viewing the details of capacity pools](#)**
You can view details of a capacity pool, such as its total capacity, used and available capacity, and member resources.
- **[Modifying capacity pools](#)**
You can modify the attributes of a capacity pool such as its name and description, and you can remove storage resources from a capacity pool.
- **[Deleting capacity pools](#)**
You can delete a capacity pool that is no longer needed.

Adding resources to capacity pools

Divide storage resources into capacity pools to later restrict provisioning or optimization to a specific set of resources. You can also use capacity pools to track the used and available capacity for block and file storage on any set of storage resources. You can add resources to an existing capacity pool, or you can create a new capacity pool.

Before you begin

You must have Administrator privileges to perform this task.

About this task

You can add storage resources to a capacity pool from resource list pages such as the Storage Systems page and the Pools page.

You can add the following resource types to a capacity pool:

- Storage systems
- Storage pools
- File systems of file storage systems

When you add a storage system to a capacity pool, any storage pool or file system that is an internal resource of that storage system is also indirectly assigned to the capacity pool.

Restriction: A resource can be assigned to only one capacity pool. When you attempt to add one or more resources to a capacity pool that are already assigned to different capacity pool, IBM Spectrum® Control displays a message that lists the resources and their capacity pool assignments. When a message is displayed, ensure that you want to change the capacity pool assignments for the listed resources.

Procedure

To add one or more resources to a capacity pool, complete these steps:

1. Open the appropriate list page for the resource you want to add. You can add an entire storage system, or you can add individual pools or file systems.
 - To add a storage system, in the menu bar, go to Storage > Block Storage Systems or Storage > File Storage Systems.
 - To add a storage pool, in the menu bar, go to Storage > Pools.
 - To add a file system, complete the following steps:
 - In the menu bar, go to Storage > File Storage Systems.
 - Select the storage system that contains the file system, then select View Details from the context-sensitive menu.
 - From the Internal Resources list, click File Systems or Network Shared Disks.
2. To add one or more resource to a capacity pool, select the resources in the list, and then select Actions > Add to Capacity Pool.
If one or more capacity pools are defined, a list of capacity pools is displayed. If no capacity pools are defined, the Create Capacity Pool window is displayed so that you can add resources to a new capacity pool.
Tip: The list of capacity pools and the Create Capacity Pool window shows capacity information. For the list of capacity pools, the Capacity column shows a bar that summarizes the used and free space in each capacity pool. In the Create Capacity Pool window, a capacity bar summarizes the used and free space of the resources that are being added to the new capacity pool. Move the mouse pointer over the bar to display a text summary and details of the capacity usage.
3. Take one of the following actions:
 - If the Create Capacity Pool window is displayed, go to step 4.
 - If a list of capacity pools is displayed, add the selected resources to an existing capacity pool or to a new capacity pool.
 - To add the selected resource or resources to an existing capacity pool, select the appropriate row in the list of capacity pools, and click Save.
 - To add the selected resource to a new capacity pool, click Create Capacity Pool. The Create Capacity Pool window is displayed.
4. In the Create Capacity Pool window, specify a name for the capacity pool.
5. Optional: Specify a capacity pool description.
Capacity pool descriptions appear in the list of capacity pools on the Capacity Pools page.
6. Optional: Specify custom tags to associate any significant information with the capacity pool.
You can specify up to three custom tags. The custom tags can be displayed in the list of capacity pools on the Capacity Pools page.
7. Save the new capacity pool by clicking OK.

Viewing the details of capacity pools

You can view details of a capacity pool, such as its total capacity, used and available capacity, and member resources.

About this task

The Capacity Pools page lists all of the capacity pools that are defined in IBM Spectrum® Control. For any capacity pool in the list, you can open a window to view more details about the capacity pool and its storage resources.

Procedure

To view details of a capacity pool, complete these steps:

1. In the menu bar, go to Advanced Analytics > Cloud Configuration, and click Work With Capacity Pools.
The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select Actions > View / Modify.
A window is displayed, showing the following details:
 - A bar summarizing the used and free space in the capacity pool in gibibytes. The bar represents the total capacity of the capacity pool. The bar is divided into sections that represent the following measurements:
 - The amount of capacity that is used for file storage
 - The amount of available capacity for file storage
 - The amount of capacity that is used for block storage
 - The amount of block storage capacity that is not used for volumes, and is not reserved by pending or scheduled provisioning tasks.

- The amount of block storage capacity that is reserved by pending or scheduled provisioning tasks
- Move the mouse pointer over the bar to display a text summary and details of the capacity usage, and a legend for interpreting the colors on the bar.
- A table that lists the storage systems, storage pools, and file systems that are in the capacity pool. To view details of any of these resources, select the resource, and then select Actions > View Properties.

Modifying capacity pools

You can modify the attributes of a capacity pool such as its name and description, and you can remove storage resources from a capacity pool.

Before you begin

You must have Administrator privileges to perform this task.

Procedure

To modify a capacity pool, complete these steps:

1. In the menu bar in the web-based GUI, go to Advanced Analytics > Cloud Configuration, and click Work With Capacity Pools.
The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select Actions > View / Modify.
A window that shows details of the capacity pool is displayed. A table lists the storage systems, storage pools, and file systems that are in the capacity pool.
3. Optional: Modify attributes of the capacity pool.
4. Optional: Remove one or more resources from the capacity pool. Removing a storage resource from a capacity pool does not delete the storage resource, but does remove its capacity from the capacity pool. To remove resources, select the appropriate resources in the list, and then select Actions > Remove.
The resources are removed from the list, but are not removed from the capacity pool until you save your changes.
5. Click Save.

Deleting capacity pools

You can delete a capacity pool that is no longer needed.

Before you begin

You must have Administrator privileges to perform this task.

About this task

When you delete a capacity pool, the storage resources that were added to the capacity pool are no longer members of any capacity pool. Deleting a capacity pool does not affect any volumes or NAS shares that were provisioned from the capacity pool. However, the volumes or NAS shares are no longer associated with a capacity pool.

Procedure

To delete a capacity pool, complete the following steps:

1. In the menu bar in the web-based GUI, go to Advanced Analytics > Cloud Configuration, and click Work With Capacity Pools.
The Capacity Pools page is displayed.
2. From the list of capacity pools, select the capacity pool, and then select Actions > Delete
3. In the confirmation window, click Delete.

Block storage: Calculating available capacity and determining the placement of volumes

For each block-storage service class, IBM Spectrum® Control calculates the amount of available capacity for volumes of that service class. When you request volumes of a particular service class, IBM Spectrum Control recommends the best location for the volumes from the available capacity.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

The Service Classes page in the GUI displays the available capacity for each block-storage service class. The available capacity is the amount of space that is available across storage pools for new volumes of the service class. Typically, not all pools are used in the available-space calculation. Some pools might be eliminated because of the consolidated status of the pool or its storage system. Other pools might be eliminated because they do not satisfy the requirements of the service class. The set of storage pools that are used to calculate the available capacity are also the candidates for provisioning by using the service class.

When you request storage volumes by using the Provision Storage wizard, you specify the capacity that is required for one or more volumes, and you specify a service class. You can also specify a capacity pool from which the volumes must be provisioned. Based on the volume size that is requested and if a capacity pool is specified, the candidates for provisioning are further refined to include only the pools that have enough available capacity and are in the specified capacity pool. From the remaining candidates for provisioning, IBM Spectrum Control creates a provisioning task that identifies the best location for the new volumes.

How available capacity for a block-storage service class is calculated

For each block-storage service class, IBM Spectrum Control calculates the available capacity for volumes of that service class. Only storage pools that are candidates for provisioning by using the service class are used to calculate the available capacity.

Identifying candidates for provisioning is a filtering process. For example, if the service class specifies that RAID 10 is required, then any storage pool that is not configured at RAID 10 is eliminated as a candidate for provisioning. IBM Spectrum Control filters the set of available storage resources against certain properties of the service class until only the storage resources that can satisfy all of the requirements remain. These remaining storage resources are the candidates for provisioning. Only certain properties of a service class are used to identify candidates for provisioning. More advanced properties contain configuration instructions, and are not used to identify candidates for provisioning.

The candidates for provisioning are identified by eliminating storage pools that cannot provide the service class, as follows:

- If a storage pool is a primordial pool or if it is formatted for count key data (CKD) volumes, it is not a candidate for provisioning.
- If a storage system has a condition of Error or Unreachable, none of its internal storage pools are candidates for provisioning. If a storage pool has a status of Error, it is not a candidate for provisioning.

Remember: *Status* is different from *condition*. Status is the status reported by a storage resource without considering the status of its internal resources. Condition is the combined status of the monitored, internal resources for the higher-level resource, and the status of the higher-level resource itself. The condition of a storage system is shown on the Storage Systems page and the condition for a pool is shown on the Pools page. The condition of a resource might not have the same value as the status, which is used to eliminate pools as candidates for provisioning.

Although the condition of a resource might not have the same value as its status, the condition is always at least as severe as the status. The condition reflects the most severe status of any internal resource or the status of the resource itself. If the condition of a resource is Normal, then its status is also Normal.

Tip: The Service Classes page in the GUI displays both the available capacity and the unavailable capacity for a service class. The Unavailable Capacity column shows the amount of storage that satisfies the requirements of the service class, but is unavailable because of the consolidated status of the storage system or its internal resources.

- If a service class allows provisioning only from certain capacity pools, any storage pool that is not in a specified capacity pool is eliminated as a candidate for provisioning.

Tip: You can see whether a service class restricts provisioning to one or more capacity pools from the Service Classes page. On the Service Classes page, show the Capacity Pools column, which is hidden by default.

- A storage pool must satisfy the requirements of every service class property in the following table to be a candidate for provisioning. Other properties of a service class are instructions for provisioning and do not affect whether the storage pool is a candidate for provisioning.

Table 1. Filtering candidates for provisioning based on service class properties

Service class property:	Candidates for provisioning:
Storage tier	If specified in the service class, storage pools that are assigned to any storage tier within the specified range are candidates for provisioning.
RAID level	If a RAID level is specified, only storage pools that are configured at the specified RAID level are candidates for provisioning. Exception: The IBM® XIV® Storage System uses a nontraditional approach to data redundancy and does not support any of conventional RAID levels. IBM Spectrum Control always considers XIV storage pools to be configured at RAID 10. Restriction: Storage pools that support multiple RAID levels are not candidates for provisioning when a service class specifies a RAID level. Storage pools that support multiple RAID levels are not candidates even when the specified RAID level is one of the RAID levels that are supported by the storage pool. Storage pools that support multiple RAID levels are only candidates for provisioning when the service class specifies that any RAID level is allowed.
Virtualization	If virtualization is required by the service class, only storage pools on IBM storage virtualizers are candidates for provisioning. The IBM storage virtualizers are SAN Volume Controller, Storwize® V7000, Storwize V7000 Unified, FlashSystem 9100, and FlashSystem V9000. If virtualization is not allowed by the service class, only storage pools that are not on IBM storage virtualizers are candidates for provisioning.
Encryption	If encryption is required by the service class, only storage pools on DS8000® storage systems that are encrypted and configured to use the encryption group that is specified in the service class are candidates for provisioning. If encryption is not allowed by the service class, only pools that are not encrypted are candidates for provisioning.
Thin provisioning	If thin provisioning is required by the service class, only storage pools that are configured for thin provisioning are candidates for provisioning. Thin provisioning is supported by the XIV, DS8000, and storage systems that run IBM Spectrum Virtualize. During provisioning, a thin-provisioned volume is created. If thin provisioning is not allowed by the service class, any storage pool that is not thin provisioned is a candidate for provisioning. In addition, any storage pool on DS8000 or a storage system that runs IBM Spectrum Virtualize is a candidate for provisioning, regardless of whether it is configured for thin provisioning. During provisioning, the volume that is created is not thin provisioned.
Compression	If compression is required by the service class, at least one I/O group of an IBM storage virtualizer must be enabled for compression. Only pools of the IBM storage virtualizer are candidates for provisioning.
Overallocation limit	If thin provisioning is required by the service class, IBM Spectrum Control identifies the storage pools that can thin provision volumes as candidates for provisioning. If an overallocation limit is specified, IBM Spectrum Control refines that list to include only the storage pools that have an overallocation percentage that does not exceed the overallocation limit. The overallocation percentage for a pool is determined by using the following measurements, as shown on the Pools page: <ul style="list-style-type: none">◦ For an XIV storage pool, the Soft Capacity divided by the Capacity.◦ For other storage system pools, the Provisioned Capacity.
Resource Tags	Resource tags further refine the candidates for provisioning. A service class can specify up to 3 resource tags. If resource tags are specified for the service class, only pools that have all the same tags are candidates for provisioning. If a pool is not tagged, any tags on the containing storage system also apply to the pool.

After the storage pools that can provide the service class are identified, IBM Spectrum Control calculates the amount of space that is available for provisioning new volumes of the service class. The calculated volume capacity is shown in the Available Capacity column on the Service Classes page.

The available-space calculation differs depending on whether the block-storage service class requires thin provisioning, as follows:

- If thin provisioning is not allowed, the available capacity is the sum of the following measurements, as shown on the Pools page:
 - For each XIV System pool, the Available Soft Capacity minus any Reserved Capacity
 - For other storage system pools, the Available Capacity minus any Reserved Capacity.
- If thin provisioning is required, the available capacity is the total of the following measurements, as shown on the Pools page, and calculations:
 - For each XIV System pool, the Available Soft Capacity minus any Reserved Capacity
 - For other storage system pools, the pool Capacity multiplied by the overallocation limit that is specified in the service class, minus the Provisioned Capacity. Any Reserved Capacity is subtracted from the result. For DS8000 storage pools, any space that is used by Track Space Efficient (TSE) volumes is subtracted

from the calculated available capacity.

How the recommended placement of storage volumes is determined

When you request one or more new volumes by using the Provision Storage wizard, the candidates for provisioning are further refined and a provisioning task is created which shows the recommended placement of the new volumes. The same filtering that was used to determine the available capacity for the service class is used, and is refined further as follows:

- A provisioning request specifies the set of storage pools from which the storage can be provisioned. This set can be all the storage pools that are known to IBM Spectrum Control, or can be limited to the storage pools in a particular capacity pool. Only the storage pools in the specified set are candidates for provisioning. This set might be more restrictive than the set used to calculate the available capacity for the service class.
- Any storage pool that cannot provide the requested volume capacity is eliminated as a candidate for provisioning.
- If virtualization and VDisk mirroring are specified by the service class, then a VDisk copy of the volume must be made in a second storage pool on the same IBM storage virtualizer. If the storage pools on the IBM storage virtualizer cannot provide the requested capacity for both the volume and its mirrored copy, the pools are eliminated as candidates for provisioning.
- If an overallocation limit is specified in the service class, only the storage pools that will not exceed the limit after the new volume or volumes are provisioned are candidates for provisioning.
Remember: Storage pools that have an overallocation percentage that does not exceed the overallocation limit, are included in the available-space calculation. However, if the overallocation percentage would exceed the overallocation limit after the requested volume space is included, the pools are not candidates for provisioning.

After the storage pools that can satisfy the provisioning request are identified, IBM Spectrum Control finds the best location for the storage. The best location for the storage is based on the unused volume capacity in the pool and performance data. Preference is first given to storage pools and systems that already contain volumes for the selected server or hypervisor. Preference is given to systems that have available performance data.

File storage: Calculating available capacity and determining the placement of shares

For each file-storage service class, IBM Spectrum® Control calculates the amount of available capacity for NAS file shares of that service class. When you request file shares of a particular service class, IBM Spectrum Control recommends the best location for the shares from the available capacity.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

For each file-storage service class, IBM Spectrum Control calculates the amount of available capacity for NAS file shares of that service class. When you request file shares of a particular service class, IBM Spectrum Control recommends the best location for the shares from the available capacity.

The Service Classes page displays the available capacity for each file-storage service class. The available capacity is the amount of space that is available across all IBM® Storwize® V7000 Unified file systems that are known to IBM Spectrum Control and that can satisfy the requirements of the service class. The set of file storage resources that are used to calculate the available capacity are also the candidates for provisioning by using the service class.

When you request a file share by using the Provision Storage pages, you specify the capacity that is required for the share, and you specify a service class. You can also specify a capacity pool from which shares must be provisioned. Based on the requested size of the share and whether a capacity pool is specified, the candidates for provisioning are further refined to include only the file storage resources that have enough available capacity and are in the specified capacity pool. From the remaining candidates for provisioning, IBM Spectrum Control creates a provisioning task that identifies the best location for the new share.

How available capacity for a file-storage service class is calculated

For each file-storage service class, IBM Spectrum Control calculates the available capacity for file shares of the service class. Only the file storage resources that are candidates for provisioning by using the service class are used to calculate the available capacity.

Identifying candidates for provisioning is a filtering process. File systems that could be candidates for provisioning if shared storage is allowed are eliminated as candidates for provisioning when dedicated storage is required. IBM Spectrum Control filters the set of available storage resources against each attribute of the service class until only the storage resources that can satisfy all of the requirements remain. These remaining storage resources are the candidates for provisioning.

The candidates for provisioning are identified by eliminating storage resources that cannot provide the service class, as follows:

- Only IBM Storwize V7000 Unified systems are supported for provisioning file shares. No other storage system type is a candidate for provisioning file shares.
- If a storage system has a consolidated status of Error or Unreachable, none of its internal NSDs or file systems are included in the available-space calculation.
- If an NSD or file system has a consolidated status of Error, it is not included in the available-space calculation.
Tip: The status of file systems is not displayed on the GUI. View the status on the NSD on which a file system resides to determine the file system status.
- If the storage volume or NSD on which a file system resides has a consolidate status of Error, the file system is not included in the available-space calculation.
Remember: *Status* is different from *condition*. Status is the status reported by a storage resource without considering the status of its internal resources. Condition is the combined status of the monitored, internal resources for the higher-level resource, and the status of the higher-level resource itself. The condition of a storage system is shown on the Storage Systems page, the condition of a volume is shown on the Volumes page, and the condition of an NSD is shown on the storage system details page. The condition might not have the same value as the status, which is used to eliminate file storage resources as candidates for provisioning.
Although the condition of a resource might not have the same value as its status, the condition is always at least as severe as the status. The condition reflects the most severe status of any internal resource or the status of the resource itself. If the condition of a resource is Normal, then its status is also Normal.
Tip: The Service Classes page displays both the available capacity and the unavailable capacity for a service class. The Unavailable Capacity column shows the amount of storage that satisfies the requirements of the service class, but is unavailable because of the consolidated status of the storage system or its internal resources.
- If a service class allows provisioning only from certain capacity pools, any storage resource that is not in a specified capacity pool is eliminated as a candidate for provisioning.
Tip: You can see whether a service class restricts provisioning to one or more capacity pools from the Service Classes page. On the Service Classes page, show the Capacity Pools column, which is hidden by default.
- A file system or NSD must satisfy the requirements of every service class property in the following table to be a candidate for provisioning. Other properties of a service class are instructions for provisioning and do not affect the storage placement determination.

Table 1. Filtering candidates for provisioning based on service class properties

Service class attribute:	Candidates for provisioning:
Shared storage / Dedicated storage	If the service class specifies dedicated storage, the share can be created only on an unused NSD. File systems and other file storage resources are eliminated as candidates for provisioning.
Resource tags	Resource tags further refine the candidates for provisioning. A service class can specify up to 3 resource tags. If resource tags are specified for the service class, only file systems and NSDs that have all the same tags are candidates for provisioning. If a file system or NSD is not tagged, any tags on the containing storage system also apply to the internal resource.

After the file storage resources that can provide the service class are identified, IBM Spectrum Control calculates the amount of space that is available for providing new files shares of the service class.

The available-space calculation differs depending on whether the file-storage service class requires dedicated storage or allows shared storage, as follows:

- If the file-storage service class requires dedicated storage, the available capacity is the total available capacity on unused NSDs within each storage system. You can verify that an NSD is unused, and view its available capacity, from the details page of a storage system. To view information about the NSDs that are associated with a storage system, click Network Shared Disks. If the NSD is unused, the File Systems column shows 0. The Available Capacity column shows the amount of space on the NSD, which is the measurement that is used in the available-space calculation for the service class.
- If the file-storage service class allows shared storage, the file system in which the NAS share is provisioned can contain other NAS shares. In this case the available capacity is the total of the following measurements:
 - The available capacity on unused NSDs within each storage system.
 - The available capacity for each file system within each IBM Storwize V7000 Unified storage system. You can view the file system available capacity from the details page of a file storage system. To view information about the file systems that are in a storage system, click File Systems. The Available Capacity column shows the amount of available capacity on each file system, which is the measurement that is used in the available-capacity calculation for the service class.

How the recommended placement of file shares is determined

When you request a file share by using the Provision Storage pages, the candidates for provisioning are further refined and a provisioning task is created which shows the recommended placement of the new share. The same filtering that was used to determine the available capacity for the service class is used, and is refined further as follows:

- A provisioning request specifies the set of file storage resources from which the storage can be provisioned. This set can be all the storage resources that are known to IBM Spectrum Control, or can be limited to the storage resources in a particular capacity pool. Only the storage resources in the specified set are candidates for provisioning. This set might be more restrictive than the set used to calculate the available capacity for the service class.
- Any storage resource that cannot provide the requested capacity for the share is eliminated as a candidate for provisioning.

The shares are created on file systems or NSDs. After the file systems and NSDs that can satisfy the provisioning request are identified, IBM Spectrum Control identifies the best location for the storage. The best location for the storage is based on the available capacity on the file system or NSD.

Changing the default host definition for provisioned storage

Use the **setdscfg** command to change the default host definition for fibre-channel ports when provisioning storage from resources that run IBM Spectrum Virtualize.

Before you begin

You must have Administrator authority to use the **setdscfg** command.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum® Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

About this task

When you provision storage from a storage system that runs IBM Spectrum Virtualize, which does not have an existing host definition for the fibre-channel ports, IBM Spectrum Control creates a new host definition during the provisioning task. The new host definition can access only one I/O group on the source device that you specified when you created the provisioning task. To allow the new host definition to access any I/O group on the source device, use the management application for the source device or use the IBM Spectrum Control **setdscfg** command.

To change the host definition to allow access to any I/O groups on the source device:

Procedure

1. Start the IBM Spectrum Control CLI by issuing the **tpctool** command.
2. Issue the **setdscfg** command with the following values for the -property and -context parameters:
 - a. Specify *Configuration.RestrictNewHostsToIogroup* as the property value.
 - b. Specify *DiskManager* as the context value and specify false as the context parameter.

For example, to allow the new host definition to access any I/O groups on the source device of the provisioning task:

```
tpctool setdscfg -url localhost:9550 -user ***** -pwd *****
-property Configuration.RestrictNewHostsToIogroup -context DiskManager false
```

Results

The value of Configuration.RestrictNewHostsToIogroup is set to false in the IBM Spectrum Control database. This removes the default restriction on the new host definition and allows access to any I/O groups on the provisioning source device.

Tip: To verify that the host definition restriction is changed, enter the following command:

```
tpctool getdscfg -url localhost:9550 -user ***** -pwd *****  
-property Configuration.RestrictNewHostsToIogroup -context DiskManager
```

- [setdscfg](#)
- [getdscfg](#)

Provisioning storage with the IBM Spectrum Control GUI

You can provision volumes or shares to servers, hypervisors, and clusters in IBM Spectrum® Control.

Before you begin

Discontinued support: Provisioning is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your provisioning needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#). The prerequisites for provisioning storage are as follows:

- You must have the required license. [Actions that are available based on role](#)
- You must have Administrator privileges or permission to use one or more of the service classes that are created for block-storage or file-service provisioning.
- You must add and probe block-storage or file-storage resources, or both.

The following restrictions apply to volume provisioning:

- If the service class that is associated with the provisioning request requires multipathing or if automatic zoning is enabled, the fabrics and switches must be managed by IBM Spectrum Control. If the fabric is not managed by IBM Spectrum Control, the fabric-related configuration options are ignored when the provisioning task is executed.
- If you provision volumes to physical servers and servers that are virtual machines, you must create separate provisioning tasks using the GUI.
- If you provision multiple servers with storage volumes, the servers must all run on the same operating system. If the servers are virtual machines, the hypervisors that manage the servers must all run on the same operating system.
- If the provisioning request involves assigning multiple volumes to a resource, the volumes are sourced from pools in the same storage system.

NAS file shares can be provisioned only from Storwize® V7000 Unified storage systems.

About this task

Specify the capacity that is required and the service class that you want to associate with the provisioning task. For shares, you must also specify at least one file access protocol. IBM Spectrum Control identifies the storage resources that can provide the storage capacity and that meet the criteria that are specified for the service class. From the set of storage resources that can provide the capacity and service class, IBM Spectrum Control identifies the best location for the storage and creates a provisioning task. Depending on the privileges that are associated with your role, you can save, execute, or schedule the provisioning task.

- [Provisioning volumes to servers](#)
You can provision volumes to one or more servers. The GUI guides you through the steps of requesting capacity and specifying requirements by using a service class.
- [Provisioning volumes to hypervisors](#)
You can provision volumes to one or more hypervisors. To provision volumes, the GUI guides you through the steps of requesting capacity and specifying requirements by using a service class.
- [Provisioning volumes to a server cluster](#)
You can assign volumes to a cluster of servers with the GUI.
- [Provisioning volumes to a hypervisor cluster](#)
You can assign volumes to a cluster of hypervisors.
- [Provisioning shares to servers](#)
You can provision a network-attached storage (NAS) file share to one or more servers. The GUI guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.
- [Provisioning shares to hypervisors](#)
You can provision a network-attached storage (NAS) file share to one or more hypervisors. The GUI guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.
- [Provisioning shares to a server cluster](#)
You can assign a network-attached storage (NAS) file share to a server cluster using the GUI.
- [Provisioning shares to a hypervisor cluster](#)
You can assign a network-attached storage (NAS) file share to a hypervisor cluster.
- [Provisioning with zone control](#)
When you provision storage, you can set a zoning policy to enable automatic zoning. When automatic zoning is enabled, IBM Spectrum Control can create zones during storage provisioning to connect a server to a storage system.

Provisioning volumes to servers

You can provision volumes to one or more servers. The GUI guides you through the steps of requesting capacity and specifying requirements by using a service class.

About this task

You can provision volumes to one or more servers. If a server is a virtual machine, the storage is provisioned to the hypervisor that is managing the server.

Procedure

1. In the menu bar, go to Servers, > Servers..
Tip: Alternatively, you can go to Advanced Analytics, > Provisioning, and click Provision to Servers.
2. Select one or more servers, and then select Actions, > Provision Storage.
3. Follow the instructions to assign the storage volumes.

Results

The GUI enables you to create a provisioning task.

Provisioning volumes to hypervisors

You can provision volumes to one or more hypervisors. To provision volumes, the GUI guides you through the steps of requesting capacity and specifying requirements by using a service class.

About this task

You can provision volumes to one or more hypervisors.

Procedure

1. Servers, > Hypervisors
Tip: Alternatively, you can go to Advanced Analytics, > Provisioning, and click Provision to Hypervisors.
2. Select one or more hypervisors, and then click Actions, > Provision Storage.
3. Follow the instructions to assign the storage volumes.

Results

The GUI creates a provisioning task.

Provisioning volumes to a server cluster

You can assign volumes to a cluster of servers with the GUI.

About this task

If the cluster is on servers that are virtual machines, the storage is assigned to the hypervisors that manage the servers. The volumes are assigned to all of the nodes in the cluster.

Procedure

1. In the menu bar, go to Servers, > Servers, and click Clusters.
Tip: Alternatively, you can go to Advanced Analytics, > Provisioning, click Provision to Servers, and click Clusters.
2. Select a cluster, and then click Actions, > Provision Storage.
3. Follow the instructions to assign the storage volumes.

Results

The GUI enabled you to create a provisioning task.

Provisioning volumes to a hypervisor cluster

You can assign volumes to a cluster of hypervisors.

About this task

The volumes are assigned to all of the nodes in the hypervisor cluster.

Procedure

1. In the menu bar, go to Servers, > Hypervisors, and click Clusters.
Tip: Alternatively, you can go to Advanced Analytics, > Provisioning, click Provision to Hypervisors and click Clusters.
2. Select a cluster, and then click Actions, > Provision Storage.
3. Follow the instructions to assign the storage volumes.

Results

The GUI created a provisioning task.

Provisioning shares to servers

You can provision a network-attached storage (NAS) file share to one or more servers. The GUI guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.

About this task

You can provision a NAS file share to one or more servers. The share is made available to the selected servers, but is not automatically mounted from the selected servers.

Procedure

1. In the menu bar, go to Servers, > Servers..
Tip: Alternatively, you can go to Advanced Analytics > Provisioning, and click Provision to Servers.
2. Select one or more servers, and then select Actions > Provision Storage.
3. Follow the instructions to assign the share.

Results

The GUI enabled you to create a provisioning task.

Provisioning shares to hypervisors

You can provision a network-attached storage (NAS) file share to one or more hypervisors. The GUI guides you through the steps of requesting capacity, specifying requirements by using a service class, and providing information about how to export the share.

About this task

You can provision a NAS file share to one or more hypervisors. The share is made available to the selected hypervisors, but is not automatically mounted from the selected hypervisors.

Procedure

1. In the menu bar, go to Servers > Hypervisors.
Tip: Alternatively, you can go to Advanced Analytics > Provisioning, and click Provision to Hypervisors.
2. Select one or more hypervisors, and then select Actions > Provision Storage.
3. Follow the instructions to assign the share.

Results

The GUI enabled you to create a provisioning task.

Provisioning shares to a server cluster

You can assign a network-attached storage (NAS) file share to a server cluster using the GUI.

About this task

You can provision a NAS file share to a server cluster. The share is made available to the servers in the cluster, but is not automatically mounted from the selected servers.

Procedure

To provision a NAS file share to a server cluster, follow these steps:

1. In the menu bar, go to Servers > Servers, and click Clusters.
Tip: Alternatively, you can go to Advanced Analytics > Provisioning, click Provision to Servers, and click Clusters.
2. Select a cluster, and then select Actions > Provision Storage.
3. Follow the instructions to assign the share.

Results

A provisioning task was created for a server cluster.

Provisioning shares to a hypervisor cluster

You can assign a network-attached storage (NAS) file share to a hypervisor cluster.

About this task

You can provision a NAS file share to a hypervisor cluster. The share is made available to the hypervisors in the cluster, but is not automatically mounted from the selected hypervisors.

Procedure

To provision a NAS file share to a hypervisor cluster, follow these steps:

1. In the menu bar, go to Servers, > Hypervisors, and click Clusters.
Tip: Alternatively, you can go to Advanced Analytics > Provisioning, click Provision to Hypervisors and click Clusters.
2. Select a cluster, and then select Actions > Provision Storage.
3. Follow the instructions to assign the share.

Results

A provisioning task was created for a hypervisor cluster.

Provisioning with zone control

When you provision storage, you can set a zoning policy to enable automatic zoning. When automatic zoning is enabled, IBM Spectrum® Control can create zones during storage provisioning to connect a server to a storage system.

About this task

Discontinued support: Provisioning is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your provisioning needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#). When IBM Spectrum Control creates a provisioning task, it identifies the best location for the new storage that satisfies the requirements of the service class. If automatic zoning is enabled, then, during provisioning, existing zones are used if the server already has connectivity to the storage system. Otherwise, one or more zones are created between a host initiator port and a controller, node, or module port.

IBM Spectrum Control provides the support for zone control and management of Brocade switches and fabrics that are managed by using Brocade Network Advisor, and Cisco switches and fabrics.

The Cisco support includes:

- Standard Cisco zone control capabilities
- Cisco Enhanced Zoning features.

Procedure

To set or modify the zoning policy, complete the following steps:

1. In the menu bar, go to Advanced Analytics > Provisioning, and click Set Zoning Policy.
2. Enable or disable automatic zoning:
 - Enable automatic zoning to allow IBM Spectrum Control to create zones during storage provisioning. The naming convention for new zones is *host_storage-system_suffix-number*. When you enable automatic zoning, you can specify the following options:
 - Zone name prefix
If you specify a zone name prefix, all zones that are automatically created by IBM Spectrum Control are prefixed with a string you specify. This prefix can help you identify which zones were automatically created, and which were created manually.
 - Make changes to the active zone set
Specify whether changes are made to the active zone set or to a new inactive zone set. If this check box is selected, one or more new zones are added to the active zone set. If the check box is cleared, a new inactive zone set is created and the new zone or zones are added to this new inactive zone set. The new inactive zone set will contain only the new zones.
 - Disable automatic zoning if you want IBM Spectrum Control to use only existing zones. If you disable automatic zoning, no changes will be made to the zoning for the fabric.
3. Click Save.

Optimizing storage tiering

To optimize the placement of volumes on storage tiers, analyze the tiering of volumes in your storage environment.

Before you begin

Discontinued support: Storage optimization is no longer supported in IBM Spectrum® Control. While the feature might still work in this release, it's recommended that you use another tool for your optimization and tiering needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#). Before you analyze tiering, complete the following tasks:

- Probe the storage virtualizers
- Collect performance data
- Set the tier level of the storage pools that you analyze and that you select as target pools

About this task

Analyzing tiering		
	1 Collect data Probe the storage virtualizers to collect information about the logical devices, such as the volumes and virtual disks, that are associated with the storage virtualizer.	
	2 Collect performance data Collect performance data to gauge the average workload activity of the volumes that you want to analyze.	
	3 Assign tier levels Assign tier levels to the pools that contain the volumes that you want to analyze and to the target pools.	
	4 Choose the source storage resources Select the resources that you want to analyze:  Servers  Hypervisors  Storage virtualizers  Storage pools  Volumes	
	5 Choose the target pools You select the target pools to place volumes that meet the tiering criteria that you specify.	
	6 Choose the performance data Choose the analysis period, the collection days, and time period that you want to use to analyze the tiering of the volumes.	
	7 Specify the tiering criteria Specify the tiering criteria and the pool saturation rates for the target pools.	

The purpose of analyzing tiering is to tier volumes based on the I/O rate or I/O density thresholds that you set. Depending on the conditions that are set, recommendations are generated.

Storage costs can be reduced, for example, by moving volumes with low workloads from higher and more expensive tiers to lower and less expensive tiers. Storage performance can be improved and used more efficiently by moving volumes with heavy workloads to the tiers that meet the workload requirements of the volumes. When you implement the recommendations, you ensure that the volumes are placed on the tiers that best match the workload requirements of the volumes.

You can optimize the placement of volumes on storage systems that run IBM Spectrum Virtualize.

Restriction: You cannot optimize the placement of volumes in data reduction pools.

Learn more: For information about the components that can be used with IBM Spectrum Control 5.2 (or later), see [IBM Spectrum Control interoperability matrix](#). Volumes can be moved to tiered storage pools on the same storage virtualizer, but volumes cannot be moved from one storage virtualizer pool to another storage virtualizer pool.

Optimizing volumes with service classes: Recommendations are not generated to move volumes that are assigned a service class to a destination pool that does not meet the requirements of the service class.

- [Investigating the capacity of tiered storage](#)
Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.
- [Setting the tier level of storage pools](#)
To analyze tiering and to balance pools, you must set the tier level of storage pools.
- [Renaming tiers](#)
Change the default names of the tiers so that they match the names of the tiers in your storage environment.
- [Analyzing tiering by servers](#)
Optimize tiering by analyzing the tiering of volumes on storage virtualizers that are connected to servers.
- [Analyzing tiering by hypervisors](#)
Optimize tiering by analyzing the tiering of volumes on storage virtualizers that are connected to hypervisors.
- [Analyzing tiering by storage virtualizers](#)
Optimize tiering by analyzing the tiering of volumes on storage virtualizers.
- [Analyzing tiering by storage pools](#)
Optimize tiering by analyzing the tiering of volumes in storage pools.
- [Analyzing tiering by volumes](#)
Optimize tiering by analyzing the tiering of volumes on storage virtualizers.
- [Tiering volumes by I/O density and I/O rate](#)
Set thresholds for tiering volumes by I/O rate or I/O density. To gain better optimization results, you can also customize the analysis of the workload activity of the pools.
- [Modifying the criteria for analyzing tiering](#)
When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

Related tasks

- [Analyzing and re-tiering volumes in pools on tier 1](#)
- [Tutorial: Collocating volumes](#)

Investigating the capacity of tiered storage

Review the capacity that is assigned to tiered storage to see whether your system needs more storage space or whether you can re-tier storage to satisfy the capacity demands for tiered storage. View the space usage trends at the current growth rate for each tier and the projected depletion date of storage space.

Before you begin

Add the pools in your storage environment to the tiers that best meet the performance requirements of your storage data.

About this task

You can view the total capacity by tier and the capacity that is consumed by the pools that are not tiered. You can also view information about recent storage growth for each tier, including the weekly and average storage growth, and the date when the storage space of the tier is expected to run out.

Procedure

1. Click Groups > Tiers.
2. Review the information that is shown about the distribution of capacity across each tier.
3. Optional: To place pools on lower or higher tiers, complete these steps:
 - a. Click the number in the Pools column.
 - b. Select the pools that you want to re-tier.
 - c. Right-click the pools, click Set Tier, and click the tier level.
4. Optional: To assign untiered pools to tiers, click Untiered Pools.

Example

If the storage space for tier 1 pools is running out, you can down-tier the volumes in the tier-1 pools to lower tiers. Alternatively, you can add more storage to meet the requirements of your tier-1 storage data.

Related tasks

- [Setting the tier level of storage pools](#)
- [Analyzing tiering by storage virtualizers](#)

- [Analyzing tiering by volumes](#)
- [Tiering volumes by I/O density and I/O rate](#)
- [Renaming tiers](#)

Setting the tier level of storage pools

To analyze tiering and to balance pools, you must set the tier level of storage pools.

About this task

Before you analyze tiering or balance pools, you must set the tier level of pools.

To analyze tiering, you must set the tier level of the source pools and the target pools. The source pools are the pools that you select for analysis. The target pools are the pools that are used to relocate volumes to lower or higher tiers of storage.

Tip: You can also set or modify the tier level of pools on the General tab of the View Properties page for the storage pool.

Child pools have the same tier level as the parent pool. So if you change the tier level of the parent pool or of the child pool, all of the pools in the parent-child relationship are set to the same tier level.

You can assign tier levels to storage systems that are not storage virtualizers. However, you cannot balance the pools or analyze the tiering of pools other than storage virtualizer pools.

Procedure

1. From the Storage menu, click Pools.
2. On the Pools page, select one or more storage pools.
3. Click Set \geq Tier, and then select a tier level.
To remove the tier level, click None.

Results

The tier level that you selected for the storage pools is shown in the Tier column.

Related tasks

- [Investigating the capacity of tiered storage](#)
- [Renaming tiers](#)

Renaming tiers

Change the default names of the tiers so that they match the names of the tiers in your storage environment.

Procedure

1. From the Groups menu, click Tiers.
2. On the Tiers page, click Rename Tiers.
3. On the Rename Tiers page, click Edit.
4. Click the names of the tiers that you want to change, type the new names, and then click Save.

Related tasks

- [Investigating the capacity of tiered storage](#)
- [Setting the tier level of storage pools](#)

Analyzing tiering by servers

Optimize tiering by analyzing the tiering of volumes on storage virtualizers that are connected to servers.

Before you begin

The servers that you select must be connected to one or more storage virtualizers.

Procedure

1. From the Servers menu, click Servers.
2. Right-click one or more servers, and then click Analyze Tiering.
Alternatively, you can click Actions \geq Analyze Tiering.

Analyzing tiering by hypervisors

Optimize tiering by analyzing the tiering of volumes on storage virtualizers that are connected to hypervisors.

Before you begin

The hypervisors that you select must be connected to one or more storage virtualizers.

Procedure

1. From the Servers menu, click Hypervisors.
2. Right-click one or more hypervisors, and then click Analyze Tiering.
Alternatively, you can click Actions_>Analyze Tiering.

Analyzing tiering by storage virtualizers

Optimize tiering by analyzing the tiering of volumes on storage virtualizers.

Procedure

1. From the Storage menu, click Storage Systems.
2. Right-click one or more storage virtualizers, and then click Analyze Tiering.
Alternatively, you can click Actions_>Analyze Tiering.

Related tasks

- [Investigating the capacity of tiered storage](#)

Analyzing tiering by storage pools

Optimize tiering by analyzing the tiering of volumes in storage pools.

Procedure

1. From the Storage_>Pools menu, click Pools.
2. Right-click one or more storage pools, and then click Analyze Tiering.
Alternatively, you can click Actions_>Analyze Tiering.

Analyzing tiering by volumes

Optimize tiering by analyzing the tiering of volumes on storage virtualizers.

Procedure

1. From the Storage menu, click Volumes.
2. Right-click one or more volumes, and then click Analyze Tiering.
Alternatively, you can click Actions_>Analyze Tiering.

Related tasks

- [Investigating the capacity of tiered storage](#)

Tiering volumes by I/O density and I/O rate

Set thresholds for tiering volumes by I/O rate or I/O density. To gain better optimization results, you can also customize the analysis of the workload activity of the pools.

Procedure

1. Select the resources that you want to analyze.
The source storage pools that are related to the resources that you selected are analyzed to determine whether they meet the workload requirements of the volumes. If the workload requirements of the volume in its current tier are not met, the volume is a candidate for relocation.

2. Select the target storage pools.
3. Include or exclude volumes in mirrored volume relationships from the analysis.
Mirrored volumes: If you want to optimize mirrored volumes, select one or more target pools that do not contain either the primary volume or the secondary volume. Primary volumes and secondary volumes cannot be placed in the same destination pool.
4. Select the period that you want to use to analyze the performance data.
You can choose the number of days, days of the week, and hours of the day for the analysis to occur.
5. Specify whether volumes that are in the same source storage pool and that are assigned to the same server or hypervisor must be placed in the same destination storage pool.
Recommendations are only generated to move one or more of the related volumes when all of the related volumes can be placed in the same destination storage pool.
Multiple host connections to the same hypervisor or server: If the volumes in the source pool that are assigned to the same hypervisor or server are assigned to different host connections, the collocation of the volumes is affected. In such cases, if volumes require optimization, the volumes that are assigned to the same host connection are kept together. To view information about the host connection for the volume, right-click the volume, select View Details, and then click the Host Connections tab.
6. Specify the type of threshold and the threshold values for moving volumes to higher and lower tiers.
7. Click Analyze.

Results

When you complete entering the information to analyze the tiering of the volumes, recommendations for tiering are shown on the Tiering Analysis page. You can create a schedule to run the analysis that is based on the analysis criteria that you entered.

Related tasks

- [Investigating the capacity of tiered storage](#)

Modifying the criteria for analyzing tiering

When a task is created after you run the tiering analysis, you can modify the criteria for tiering the volumes.

Before you begin

You must run the tiering analysis to create a tiering analysis task.

About this task

The criteria that you modify are used to analyze the source volumes and target volumes that were selected when you created the tiering analysis.

Procedure

1. From the Home menu, click Tasks.
2. Right-click a tiering analysis task, and click Edit.
3. Modify the tiering thresholds, and then click Analyze.

Results

The source volumes are analyzed to determine whether one or more of the volumes require re-tiering based on the criteria that you specified. If one or more of the volumes require re-tiering, recommendations are generated to move the volumes to the destination pools that meet the criteria that you specified.

Optimizing storage pools

Analyze the activity of pools and resolve performance hot spots by redistributing volumes across each storage tier.

Before you begin

Discontinued support: Storage optimization is no longer supported in IBM Spectrum® Control. While the feature might still work in this release, it's recommended that you use another tool for your optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

Before you balance pools, complete the following tasks:

- Probe the storage virtualizer. You must probe the storage virtualizer to collect information about the logical devices, such as volumes and virtual disks, that are associated with the storage virtualizer.
- Collect performance data. To gauge the average activity level of the volumes and pools, you must collect sufficient performance data for the storage virtualizer.
- Assign a tier level to the pools that you want to balance.

About this task

You can balance the workload of volumes in pools on the same tier on storage systems that run IBM Spectrum Virtualize.

Restriction: You cannot balance the workload of volumes in data reduction pools.

Learn more: For information about the components that can be used with IBM Spectrum Control 5.2 (or later), see [IBM Spectrum Control interoperability matrix](#).

To balance pools, you must select two or more pools on the same tier level and on the same storage virtualizer. For example, you select three tier 1 pools, and four tier 3 pools. The pools on tier 1 are analyzed and balanced, and the pools on tier 3 are analyzed and balanced.

Optimizing volumes with service classes: Recommendations are not generated to move volumes that are assigned a service class to a destination pool that does not meet the requirements of the service class.

- **Balancing pools**

You balance pools to distribute the workload of volumes across pools on the same tier and on the same storage virtualizer. To improve the performance of the pools, recommendations are generated to move volumes to other pools on the same tier and on the same storage virtualizer.

- **Modifying the criteria for balancing pools**

When a task is created after you run the analysis for balancing pools, you can modify the criteria that determine which pools are balanced.

- **Criteria for identifying the pools that require balancing**

To identify the pools that require balancing, performance data is collected to determine the activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

Related information

- [Setting the tier level of storage pools](#)

Balancing pools

You balance pools to distribute the workload of volumes across pools on the same tier and on the same storage virtualizer. To improve the performance of the pools, recommendations are generated to move volumes to other pools on the same tier and on the same storage virtualizer.

Before you begin

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum® Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

Procedure

1. From the Storage menu, click Pools.
2. Right-click two or more pools on the same tier level and on the same storage virtualizer, and then click Balance Pools.
Tip: If you want to analyze pools that contain volumes in mirrored volume relationships, select one or more pools that do not contain either the primary volume or the secondary volume. The primary volume and the secondary volume cannot be placed in the same destination pool.
3. Choose whether to analyze volumes in mirrored volume relationships.
These options are only available if you select one or more volumes with mirrored volumes.
4. Select the period that you want to use to analyze the performance data.
5. Specify the placement of optimized volumes in capacity pools.
These options are only available if you select one or more volumes in capacity pools.
6. Specify whether volumes that are in the same source storage pool and that are assigned to the same server or hypervisor must be placed in the same destination storage pool.
7. Click Analyze.

Results

When you finish entering the information for balancing the pools, a task is created that shows the recommendations on the Balance Analysis page. You can implement the recommended changes, or you can create a schedule to implement the changes later.

Analyzing volumes with service classes: Recommendations are not generated to move volumes that are assigned a service class unless the destination pool can meet the requirements of the service class.

Modifying the criteria for balancing pools

When a task is created after you run the analysis for balancing pools, you can modify the criteria that determine which pools are balanced.

Before you begin

To create a balance analysis task, you must run the analysis for balancing the pools.

About this task

The criteria that you modify are used to analyze the storage pools that you selected when you entered the criteria for balancing the pools.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum® Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

Procedure

1. From the Home menu, click Tasks.
2. Right-click a balance analysis task, and click Edit.

3. Make your changes, and then click Analyze.

Results

Based on the criteria that you specified, the pools are analyzed to determine whether the pools require balancing. If one or more pools require balancing, recommendations are generated to redistribute the volumes to pools on the same tier of storage.

Criteria for identifying the pools that require balancing

To identify the pools that require balancing, performance data is collected to determine the activity of pools on the same tier. The performance capability of pools is preserved by setting an activity limit for each tier of storage.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum® Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning and optimization needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

To determine the activity of pools on the same tier and on the same back-end storage system, the current activity level of the pool is calculated. On the Pools page, the activity level of the pool is shown in the Activity column.

The following formulas are used to calculate the activity level of pools.

Table 1. Formulas for calculating the activity level of pools

Value calculated	Formula that is used in the calculation
Activity level of the pool	$[\text{Read I/O Rate} \times (1 - \text{Read I/O Cache Hit \%}) + \text{Write I/O Rate}] \div \text{Total Pool Capacity}$
Activity level of the pool on XIV® systems	$(\text{Total I/O Rate} \div \text{Total Capacity})$

Reclaiming storage

Use the recommendations to reclaim capacity before you plan new capacity purchases.

Before you begin

To determine which volumes can be reclaimed, capacity and performance data must be available for the reclamation analysis. To run the reclamation analysis, capacity data is collected for the previous day, and a daily aggregation of the performance data is collected for the previous 14 days.

Tip: Depending on the I/O activity of the volumes in your storage environment, you can shorten or lengthen the default period of 14 days.

Storage systems that support reclamation: To find out, go to [IBM Spectrum Control interoperability matrix](#), and then click the release number in the Storage column.

About this task

Got an alert that your storage system is running out of capacity? Or, did you click Storage > Pools, had a look at the values in the Zero Capacity column and saw that you were running out of capacity?

Instead of purchasing more capacity, click Advanced Analytics > Reclamation to see how much capacity you can reclaim.

The reclamation analysis is run daily and if sufficient data is collected recommendations are generated to reclaim the volumes that meet either one of the following criteria:

- The volume isn't assigned to a server.
- I/O activity was not detected in the data that was collected for the volume.

For example, the reclamation analysis is run, and it is determined that a volume isn't assigned to a server. The volume is identified as reclaimable even if I/O activity is detected for the volume. Alternatively, the reclamation analysis is run, and no I/O activity is detected for the volume. The volume is identified as reclaimable even if the volume is assigned to a server.

The reclamation analysis also detects if volumes are replica volumes, VDisk mirrored volumes, or FlashCopy® volumes. The same criteria are used to identify whether the volumes are reclaimable, but the process varies depending on the type of copy volume:

Replica volumes

The source volumes of volumes with replicas, such as volumes that use Metro Mirror or Global Mirror copy services, are analyzed to determine whether the volumes are reclaimable.

VDisk mirrored volumes

Both copies of the VDisk mirrored volume are analyzed to determine whether the volumes are reclaimable.

FlashCopy volumes

The target volumes of volumes in FlashCopy relationships are analyzed if the source volumes are identified as reclaimable.

After the analysis is run, you see the total capacity that can be reclaimed, which is broken down by tier, and by volume. (The reclaimable capacity for volumes in storage systems that use data reduction technologies isn't shown. See the restrictions below.) In the table, you get a list of the volumes, the volume's capacity, and the information that you need to decide whether you want to decommission the volumes.

Tip: See volumes that are identified as reclaimable, such as volumes that are used to back up data or volumes that you recently assigned to a new application? If you don't want to include these volumes in the reclamation analysis, right-click the volumes and click Exclude from Analysis.

Before you reclaim space:

It's a good practice to verify that the identified volumes are available for reclamation. Before you delete or reclaim space that was identified by IBM® Storage Insights, keep in mind the following additional considerations:

- For volumes on IBM Spectrum Accelerate and CKD volumes on DS8000®, the volumes are identified as reclaimable based on I/O activity, because information about the assignment of volumes to servers is not available.
- For non-IBM storage systems, volumes are identified as reclaimable based on I/O activity because information about their existing server assignments, replication relationships, and snapshot targets might not be available for the storage system. Check with your storage administrators to ensure that the identified volumes are available for reclamation and are not part of these other configurations.
- Volumes in storage systems that use data reduction technologies are identified as reclaimable, but the actual physical capacity of the volumes can't be determined. For example, when data is deduplicated, multiple volumes can share identical blocks of data so we don't know the actual capacity of each volume. Because we can't determine the capacity of the individual volumes, the total reclaimable capacity for each volume is displayed as zero. This applies to:
 - Volumes in IBM storage systems that support data reduction pools or volumes in storage systems that run IBM Spectrum Virtualize.
 - Volumes in IBM FlashSystem® A9000 storage systems.

Private volumes in Dell EMC storage systems are excluded from the reclamation analysis.

- **[Viewing storage reclamation](#)**
Use the storage reclamation analysis to see information about the reclaimable capacity in your data center. You can see the savings that can be made by reclaiming capacity for tiered and non-tiered storage and view a list of the reclaimable volumes.
- **[Viewing storage reclamation by storage system](#)**
Use the reclamation view of storage systems for detailed information about the reclaimable capacity in your data center. You can see the savings that can be made by reclaiming capacity for tiered and non-tiered storage and view information about the storage systems that contain the reclaimable volumes.
- **[Excluding volumes from reclamation analysis](#)**
You can select volumes to not include them in the analysis for reclamation recommendations. The volumes are excluded from the reclamation charts.

Viewing storage reclamation

Use the storage reclamation analysis to see information about the reclaimable capacity in your data center. You can see the savings that can be made by reclaiming capacity for tiered and non-tiered storage and view a list of the reclaimable volumes.

Before you begin

Before you view recommendations for reclaiming storage, complete the following tasks:

- Add the storage systems for monitoring and schedule a storage systems probe.
- If your storage environment is configured for storage tiers, ensure that the amount of available capacity on tier one is minimal.

By default, the period for analyzing performance data, which is used for generating reclamation recommendations, is set to 14 days. For a more thorough analysis, you can change this period.

About this task

View information about volumes that are analyzed and recommended for reclamation.

Procedure

1. From the Advanced Analytics menu, click Reclamation.
2. On the View by Reclaimable Capacity page, view the charts and tables for information about reclaimable storage space, volumes that are recommended for reclamation, and volumes that are excluded from analysis.

Results

On the reclamation donut chart, you can view the amount of storage space that is used and the amount of storage space that can be reclaimed. You can also see an estimate of the storage space that can be saved when the volumes that are listed in the table are reclaimed.

On the reclamation by tier bar charts, view the amount of storage space that can be reclaimed on each tier of storage that is defined in your data center.

On the Recommendations tab, view the volumes that are identified as potential candidates for reclamation.

On the Excluded tab, view the volumes that are excluded from the analysis for reclamation recommendations.

Related tasks

- [Adding resources](#)
- [Optimizing storage tiering](#)
- [Modifying the period for analyzing performance data](#)

Viewing storage reclamation by storage system

Use the reclamation view of storage systems for detailed information about the reclaimable capacity in your data center. You can see the savings that can be made by reclaiming capacity for tiered and non-tiered storage and view information about the storage systems that contain the reclaimable volumes.

Before you begin

Before you view recommendations for reclaiming storage, complete the following tasks:

- Add the storage systems for monitoring and schedule a storage systems probe.
- If your storage environment is configured for storage tiers, ensure that the amount of unused capacity on tier one is minimal.

By default, the period for analyzing performance data, which is used for generating reclamation recommendations, is set to 14 days. For a more thorough analysis, you can change this period.

About this task

View reclamation information about storage systems that contain volumes that are analyzed and recommended for reclamation.

Procedure

1. From the Advanced Analytics menu, click Reclamation.
2. On the View by Storage Systems page, view the charts and table for information about reclaimable storage space and for information about the storage systems that contain the reclaimable volumes.

Results

On the reclamation donut chart, you can view the amount of capacity that can be reclaimed and the amount of capacity that can't be reclaimed. If you defined storage tiers, you can view the amount of storage that can be saved for each tier. You can also see an estimate of the storage space that can be saved when the volumes of the storage systems that are listed in the table are reclaimed.

On the reclamation bar chart, view the amount of storage space that can be reclaimed on each storage system in your data center.

In the table, view space, capacity, and reclamation information about each storage system.

Related tasks

- [Adding resources](#)
- [Optimizing storage tiering](#)
- [Modifying the period for analyzing performance data](#)

Excluding volumes from reclamation analysis

You can select volumes to not include them in the analysis for reclamation recommendations. The volumes are excluded from the reclamation charts.

About this task

On the View by Reclaimable Capacity page, you can select a volume or multiple volumes for exclusion from the analysis for reclamation recommendations.

Procedure

1. From the Advanced Analytics menu, click Reclamation.
2. On the View by Reclaimable Capacity page, to exclude a volume from the analysis, right-click the volume on the Recommendations table, and select Exclude from Analysis.

Results

The volume is removed from the Recommendations table and added to the Excluded table, and the charts are refreshed.

Tip: To include a volume in the analysis, on the Excluded table, right-click the volume and select Include in Analysis. The volume is removed from the Excluded table and added to the Recommendations table, and the charts are refreshed.

Transforming and migrating volumes

You can transform fully allocated volumes to compressed or thin-provisioned volumes, or transform compressed or thin-provisioned volumes to fully allocated volumes. You can move volumes to other pools or to pools that are enabled for Easy Tier®.

Before you begin

Discontinued support: Storage optimization is no longer supported in IBM Spectrum® Control. While this feature might still work in this release, it's recommended that you use another tool for your volume transformation needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#). Before you transform storage, you must probe the storage virtualizers. The probe collects information about the logical devices, such as volumes and virtual disks, that are associated with the storage virtualizer.

About this task

You can convert or move volumes in pools on storage systems that run IBM Spectrum Virtualize.

For volumes in mirrored volume relationships, you can convert the primary volume, the secondary volume, or both volumes.

Restrictions: The restrictions for converting or moving volumes are as follows:

- You cannot convert or move volumes in data reduction pools.
- To move volumes, the destination pool must be on the same storage virtualizer. You cannot move a volume from a pool on one storage virtualizer to a pool on another storage virtualizer.
- To convert or move volumes, image mode volumes must be converted to managed mode volumes.
- To convert fully allocated volumes to compressed volumes, you must have the IBM® Real-time Compression license.

Procedure

1. In the menu bar in the web-based GUI, go to Storage > Volumes.
2. Click one or more storage volumes, and then select Transform Storage from the context-sensitive menu.
Alternatively, you can click Actions > Transform Storage.

Results

When you complete entering criteria for transforming storage, the Transform Plan page is opened. A task is created which you can run to implement the recommendations.

Related tasks

- [Modifying the period for analyzing performance data](#)

Modifying the period for analyzing performance data

To change the default period for analyzing performance data in storage virtualizer pools, issue the **setdscfg** command. You can extend or shorten the default period for a more accurate analysis.

Before you begin

You must have Administrator authority to use the **setdscfg** command.

About this task

By default, the period for analyzing performance data is set to 14 days. When you change the number of days, the period for analyzing performance data is modified for the advanced analytics about storage, that is, tiering, balancing, transforming, and reclamation.

When you transform storage, the performance data that is collected about the pools on storage virtualizers is analyzed. The period that is set for analyzing performance data is used to calculate the average workload activity of the volumes that are selected for analysis. To provide a more accurate analysis of the workload activity of the volumes in your storage environment, you can extend or shorten the period.

For reclamation, the performance data that is collected for the analysis period about all block storage pools is analyzed, regardless of whether the storage is virtualized.

To change the analysis period, complete the following steps:

Procedure

1. To start the IBM Spectrum® Control CLI, issue the **tpctool** command.
2. Issue the **setdscfg** command, and then set the following values by using the **-property** and the **-context** parameters:
 - a. Specify *tiering_pm_days* as the property value.
 - b. Specify *tiering* as the context value, and then specify the number of days for the analysis period.

For example, to change the analysis period from the default value of 14 days to 28 days, issue the following command:

```
tpctool setdscfg -url localhost:9550 -user ***** -pwd *****  
-property tiering_pm_days -context tiering 28
```

Results

The default value is changed to the number of days that you specified, for example 28.

Tip: To verify that the analysis period is changed, you enter the following command:

```
tpctool getdscfg -url localhost:9550 -user ***** -pwd ***** -context tiering
```

- [setdscfg command](#)
- [getdscfg command](#)

Modifying the properties of resources

Add or change the properties for individual resources or for multiple resources. You can use the properties to filter or sort the resources in the GUI or in an external application if the data is shared or exported.

Before you begin

To add or modify the properties of resources, you must have Administrator privileges.

About this task

For Servers, storage systems, hypervisors, switches, and fabrics, you can modify the properties of a single resource or multiple resources. You can modify the properties of a single pool, but not the properties of multiple pools.

Modifying the properties of multiple resources

1. From the menu, click the type of resource that you want to edit. For example, if you want to edit the properties for servers, click **Resources > Servers**.
2. To select the resources, click Ctrl + click or Shift + click.
3. Right-click the selected resources and click **View Properties**.
The properties can have the following values:
 - If the property is blank, it means that the values were never assigned or the resources that were selected have different values assigned to that property. For example, the selected servers have different locations.
 - If the selected resources have the same value for a property, the value is shown.
4. In the Properties notebook, click **Edit**.
5. Modify the properties and click **Save**.

Modifying the properties of a single resource

1. From the menu, click the type of resource that you want to edit. For example, if you want to edit the properties for servers, click **Resources > Servers**.
 2. Right-click the resource that you want to modify and click **View Properties**.
 3. In the Properties notebook, click **Edit**.
 4. Modify the properties and click **Save**.
- **Properties of resources that can be modified**
You can modify the properties, such as the name or location, of the resources that are monitored by IBM Spectrum Control.

Properties of resources that can be modified

You can modify the properties, such as the name or location, of the resources that are monitored by IBM Spectrum® Control.

The following table lists the properties that you can modify for each type of resource:

Resource	Properties that can be modified
Storage systems	<ul style="list-style-type: none">• Name (label)• Location• Custom tags (3)
Pools	<ul style="list-style-type: none">• Tier• Back-end storage system type• Back-end storage RAID level• Back-end storage disk type• Back-end storage disks <p>Tip: You can edit the back-end values if the back-end storage system for the related storage virtualizer was not probed and the values for the pool are unknown. If you manually enter values for these properties, IBM Spectrum Control uses the values to help calculate the approximate read I/O capability of the pool.</p>
Servers	<p>Servers with Storage Resource agents</p> <ul style="list-style-type: none">• Location• Custom tags (3)• Agent trace• Trace level• Number of trace files• Trace file size• Run scripts on agent <p>Servers without Storage Resource agents (agentless)</p> <p>You can modify the following properties for agentless servers:</p> <ul style="list-style-type: none">• Name (label)• OS type• IP address• Location• Custom tags (3)
Hypervisors	<ul style="list-style-type: none">• Location• Custom tags (3)
Switches	<ul style="list-style-type: none">• Name (label)• Location• Custom tags (3)

Resource	Properties that can be modified
Fabrics	<ul style="list-style-type: none"> Name (label) Location Custom tags (3)

Opening the management GUI for a resource

Resources such as storage systems and switches can have their own management GUIs or element managers. In IBM Spectrum® Control, you can open the start page in these GUIs from resource list pages, such as the Storage Systems page and Switches page.

About this task

The action for opening a management GUI is available for storage systems and switches under the following conditions:

- The resource is monitored by IBM Spectrum Control.
- The resource has a management GUI that is web-based.

Procedure

- In the menu bar in the GUI, go to a type of top-level resource.
For example, if you want open the GUI for a block storage system, go to Storage > Block Storage Systems. If you want open the GUI for a switch, go to Network > Switches.
- Right-click the resource and select the action to open the GUI. For storage systems, select Launch Storage System GUI. For a switch, select Open Switch GUI.
The start page for the management GUI is opened in a separate web browser window.

Opening the DS Storage Manager GUI for a DS8000

About this task

If the firmware level for a DS8000® is earlier than R6.2, you must complete extra configuration steps before you can open its management GUI. To determine the firmware version of a DS8000 and enable the ability to open its management GUI, complete the following steps:

Procedure

- Start IBM Spectrum Control.
- In the menu bar, go to Storage > Block Storage Systems.
- Right-click the DS8000 and select View Properties.
- Check the version in the Firmware field. If the firmware version is earlier than R6.2, continue to the next step. If the version is R6.2 or later, you do not need to complete this procedure.
- Download and start the PuTTY utility. PuTTY is a free implementation of Telnet and SSH for Windows and UNIX. You can download PuTTY from <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>.
- On the Session page in the PuTTY tool, complete the following actions:
 - In the Host Name (or IP address field), enter the IP address of the Hardware Management Console (HMC) for the DS8000. You can determine this IP address by viewing the properties of the DS8000 in the IBM Spectrum Control GUI.
 - Select SSH for a Connection type.
 - Click Open.
- Enter the user name and password for logging in to the HMC.
- Run the following command: `touch /opt/esshmc/data/config/ntvbrsen`
- In the GUI, right-click the DS8000 and select Launch Storage System GUI.

Exporting information to a file

You can save information about resources, tasks, or alerts to a PDF, CSV, or HTML file. Information that you can export from the GUI includes all the values that are being shown in the columns for a list of resources, tasks, or alerts.

Before you begin

The information that you export to a file is organized according to the sorting, filtering, and column order that is defined for a list. Before you export information, complete the following tasks to configure a list:

- Change the order of columns in the list as you want them to appear in the generated file.
- Hide columns in the list that you do not want to include in the generated file. Information in hidden columns is not exported.
- Sort the rows in the list as you want them to appear in the generated file. Each list in the GUI has a column or set of columns that determine the order of its rows.
- Filter the list to show only the resources, tasks, or alerts that you want to export. Use filtering to limit the rows that are shown in a list based on value in a specific column or set of default columns.

About this task

You can export information that is shown on the following pages:

Resource list and resource details pages

You can export information about the top-level resources, internal resources, or related resources that are shown on resource list pages and resource details pages. For top-level resources, you can also export information about related tasks or alerts. For example, you can export information about monitored switches on the Switches page. When you select the export action, the generated file includes the column values for each monitored switch in the list. In the Alerts tab on the Switches page, you can export information about all the alerts that are related to switches.

Alerts page

On the Home > Alerts page, you can export information about all the alerts that were detected for monitored resources. Information that you can export includes all the values that are being shown in the columns for the list of alerts.

Tasks page

On the Home > Tasks page, you can export information about all the tasks that are used to optimize resources and provision storage. Information that you can export includes all the values that are being shown in the columns for the list of tasks.

Procedure

1. In the GUI, go to a list of resources, tasks, or alerts for which you want to export information.
2. Right click any row in the table and select Export > `<file_format>`, where `<file_format>` represents the format that you want to use for a file. The following formats are available:

CSV (comma-separated values)

A CSV file is a file that contains comma-delimited values and can be viewed with a text editor or imported into a spreadsheet application. The information in a CSV file has the following format:

```
"Column_name", "Column_name", "Column_name", "Column_name", "Column_name", "Column_name"
"data", "data", "data", "data", "data", "data"
```

where `Column_name` represents the name of a column in a table and `data` represents the data that is associated with a column.

PDF

You can view a PDF file with Acrobat reader. The information that you export to a PDF file is formatted into a table.

Tip: When you export a table of data that contains many columns, the rows in that table might span multiple pages in the resulting PDF. For example, if a table contains 20 columns, then the row for a specific resource might be shown on more than one page (10 columns on the first page and 10 columns on the second page).

To reduce the number of pages that a table spans, before you select the export action, hide the columns that you do not want to include in the PDF. Continue hiding columns until the table no longer spans multiple pages.

HTML

You can view an HTML file with a web browser. The information that you export to an HTML file is formatted into a table.

3. Optional: Depending on how your web browser is configured, you can specify the name and location of the generated file.
4. Click the save option to export the information to a file.

Related concepts

- [Viewing information about resources](#)
- [Customizing lists of resources, tasks, and alerts](#)

Customizing lists

Customize lists to focus on the information that is important to you. You can filter information, sort rows, and show, hide, and reorder columns in the table views for resources, tasks, alerts, and other objects.

Filtering lists

You can filter the items that are shown in a list. Use filtering when you want to limit the rows that are shown in a list based on values in a specific column or set of columns. For example, when viewing alerts, you can filter the list so that only alerts with a value of Warning in the Severity column are shown.

Sorting lists

Each list of resources, tasks, alerts, or other objects in the GUI has a column or set of columns that determines the order of its rows. You can sort the columns in a list to organize the rows according to your requirements.


Showing, hiding, and reordering columns in lists

Each list in the GUI has a set of columns that contain information about resources, tasks, alerts, or other objects. You can change the order of columns, show columns that are hidden, and hide columns that are shown in these lists.

Filtering lists

You can filter the items that are shown in a list. Use filtering when you want to limit the rows that are shown in a list based on values in a specific column or set of columns. For example, when viewing alerts, you can filter the list so that only alerts with a value of Warning in the Severity column are shown.

Procedure

1. In the web-based GUI, go to a list of resources, jobs, or alerts that you want to filter.
2. Click the Filter icon  to determine how to filter the list. The following options are available:

Filter by Default Columns

Select this option to apply the filter to the set of the default columns in a list. Each list contains a set of default columns on which you can filter.

For example, in the lists of alerts on the Home > Alerts page, select this option to apply filter text to the following columns at the same time: **Condition**, **Severity**, **Alert Category**, **Resource**,

Internal Resource, Total Occurrences, Alert Name.

Filter by a Specific Column

Select this option to apply filter text to a specific column in a list. For example, on the Home_>Alerts page, select the Condition column if you want to apply filter text to values in that column only.

Restriction: Some columns in lists cannot be used for filtering and are not shown under the **Filter by a Specific Column** option.


3. In Filter text box , type the text for filtering a list.

You can type letters, whole words, partial words, and numbers. When filtering on certain columns, you can select from a list of values rather than typing in filter text. For example, when filtering on the Status column, you can select one of the following values as the filter text: **Error**, **Unreachable**, **Warning**, **Normal**, **Unknown**.

4. Press Enter to apply the filter to the list.

Only rows with values that match or partially match the filter text are shown.

For example, on the Home_>Alerts page, if you select to filter on the Severity column and type `critical` as the filter text, only alerts that have a severity of `critical` and `critical - acknowledged` are shown in the list.

Tip: The name of a column to which a filter applies is shown next to the Filter icon .


5. Optional: To remove the filtering for a list, click Reset.

Sorting lists

Each list of resources, tasks, alerts, or other objects in the GUI has a column or set of columns that determines the order of its rows. You can sort the columns in a list to organize the rows according to your requirements.

About this task

For example, the list of ports on a fabric is ordered by switch name, blade slot, and port number, in that order. You can change the sort order of these columns to organize a list of resources according to your requirements.

Tip: A visual indicator  is shown next to the primary column that determines the sort order of a list.

Procedure

1. In the GUI, go to a list of resources, tasks, alerts, or other objects that you want to sort.
2. Complete the following tasks to sort the rows in the list:
 - To change whether the rows in a list are sorted by ascending or descending order, click the name of the primary column in the heading row. Click the name of the column again to reverse the sort order.
 - To change the column that determines the sort order of the list, click the name of that column. For example, click the Name column to sort the list of items alphabetically, or click the Status column to sort the list according to the status of items.

Tip: When you click a column other than the primary sort column, that column becomes the new primary sort column. However, the previous primary column maintains its sort order (ascending or descending).

For example, if a list of resources is ordered by the Name column in ascending order and you click the Status column, the rows are sorted according to the statuses of resources. At the same time, the resources for each type of status are still sorted in ascending order by the Name column. There is a maximum of three columns that can be sorted upon at the same time.
3. Optional: To reset the order of the rows for a column to the default setting, right-click the heading row in the list and select Restore Default View.

Example

On the Switches page, you can use the sorting function to group virtual or logical switches with their physical switch. First, show the Virtual and IP Address columns. Click the Virtual column to sort it in ascending alphabetical order. Then, click the IP Address column.

Showing, hiding, and reordering columns in lists

Each list in the GUI has a set of columns that contain information about resources, tasks, alerts, or other objects. You can change the order of columns, show columns that are hidden, and hide columns that are shown in these lists.

Procedure

1. In the GUI, go to a list of resources, task, alerts, or other objects that you want to customize.
2. Complete the following tasks to customize the columns in the list:
 - To show or hide the columns in a list, right-click any of the column headings and select or clear the check box next to a column name.
 - To change the order of columns, click a column heading and drag it to a new position in the list.

Tip: The changes that you make to column settings are saved from session to session.
3. Optional: To reset the order and display of columns to the default setting, right-click any of the column headings and select Restore Default View.

Managing tasks

Tasks are used to provision storage and optimize resources in your storage environment. Use the Tasks page to manage all the tasks that are used by IBM Spectrum® Control to provision storage and optimize resources. Use resource list and resource details pages to manage the tasks for specific resources and resource types.

Discontinued support: Cloud configuration, provisioning, and optimization are no longer supported in IBM Spectrum Control. While these features might still work in this release, it's recommended that you use another tool for your provisioning, optimization, and volume transformation needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

Tasks are created when you take any of the following actions in the GUI:

- You can provision storage to assign storage to servers or hypervisors.
- You can analyze tiering to move volumes to higher or lower tiers.
- You can transform storage to move or convert volumes.
- You can balance pools to distribute the workload of volumes across pools on the same tier.

Use the following actions to manage tasks:

- View the status of tasks in a central location.
- Identify all the tasks that are related to a specific resource or a type of resource.
- Identify all the tasks that are encountering problems when the tasks are running.
- Administer the tasks. The actions that are available depend on the type of task that you select. For example, you can run or rename a tiering-analysis task, or export provisioning task details to a text file.
- Change the number of task runs that are displayed.

Tip: A *task run* is a single invocation of a task.

- **[Viewing tasks](#)**
Use the Tasks page to view all the tasks that are used by IBM Spectrum Control to provision storage and optimize resources. Use resource list and resource details pages to view the tasks for specific resources and resource types.
- **[Managing tasks for provisioning](#)**
Provisioning tasks are used to create storage volumes and assign the volumes to servers or hypervisors. Provisioning tasks are also used to create network-attached storage (NAS) file shares. Use the Tasks page in the GUI to manage all the tasks that are used by IBM Spectrum Control to provision storage. Use resource list and resource details pages to manage the provisioning tasks for specific resources and resource types.
- **[Managing tasks for tiering storage, balancing pools, and transforming storage](#)**
To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the Tasks page to manage all the tasks that are used to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.
- **[Renaming tasks](#)**
You can change the name of a task by using the task detail page. For example, use the Provisioning page to change the name of a provisioning task. A task name is automatically generated when the task is created. If the original name is ambiguous or is not clear enough, you can change the task name to a unique task name that you provide.
- **[Deleting tasks](#)**
Use the Tasks page to delete tasks, such as provisioning, tiering analysis, balance analysis, transform plan, and analysis-execution tasks.

Viewing tasks

Use the Tasks page to view all the tasks that are used by IBM Spectrum® Control to provision storage and optimize resources. Use resource list and resource details pages to view the tasks for specific resources and resource types.


- **[Viewing all tasks](#)**
Use the Tasks page as a central location to view the tasks that IBM Spectrum Control is using to provision storage and optimize resources. You can view all the tasks for resources that are monitored by IBM Spectrum Control.
- **[Viewing pending tasks](#)**
Some tasks might be pending because they can be run or scheduled only by an administrator. If tasks are pending, the number of pending tasks is displayed in the title bar of the GUI.
- **[Viewing tasks for a type of resource](#)**
Use resource list pages to view only the tasks for a specific type of top-level resource. Each resource list page has a Tasks tab that lists only the tasks for that type of resource. For example, you can use the list page for storage systems to view only the tasks that are associated with storage systems.
- **[Viewing tasks for a specific resource](#)**
Use resource detail pages to view only the tasks for a specific monitored resource. Each resource detail page has a Tasks link that lists only the tasks for the associated storage system, server, hypervisor, switch, or fabric.
- **[Viewing task details](#)**
Use task detail pages to view detailed information about the tasks that IBM Spectrum Control uses to optimize resources and provision storage. You can use this information to troubleshoot the reason why a task failed, analyze details of a task before you implement task recommendations, or schedule future task runs.
- **[Viewing task logs](#)**
Task logs include detailed information about the status, actions, and progress of a task. Status information includes informational, warning, and error messages that are related to each action in a task that is taken during processing. You can use this information to troubleshoot any errors that might occur when a task is run.
- **[Setting the number of task runs that are displayed](#)**
You can change the number of task runs that are shown in the GUI by configuring the settings on the History Retention page. By default, the information for the last five runs of a task is shown.

Viewing all tasks

Use the Tasks page as a central location to view the tasks that IBM Spectrum® Control is using to provision storage and optimize resources. You can view all the tasks for resources that are monitored by IBM Spectrum Control.

Procedure

1. In the menu bar in the web-based GUI, go to Home > Tasks.
The information about tasks is organized into columns. These columns include information about the status of tasks, the date and time of the last run for tasks, and the schedule for tasks.

2. Optional: View the status icons at the top of the page to view a summary of task statuses.
This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.
Tip: To view descriptions of the icons and columns on the Tasks page, click the Help icon  in the upper-right corner of the page.

Viewing pending tasks

Some tasks might be pending because they can be run or scheduled only by an administrator. If tasks are pending, the number of pending tasks is displayed in the title bar of the GUI.

Before you begin

You must be assigned to the Administrator role to view the number of pending tasks.

About this task

Pending tasks are tasks that were not run and are not scheduled to run. Pending tasks might be created by a user who is not authorized to run or schedule the task. For example, users who are assigned to the Monitor or External Application role can have permission to create provisioning tasks by using a service class. However, if the service class specifies that administrator approval is required, these users can only create the provisioning task. The provisioning task is pending until an administrator runs or schedules it.

Procedure

To view pending tasks, complete the following steps:



1. In the title bar of the GUI, look for the pending task count. If no number appears in the title bar, there are no pending tasks.
2. To open a Tasks page that shows only the pending tasks, click the number of pending tasks in the title bar.

Viewing tasks for a type of resource

Use resource list pages to view only the tasks for a specific type of top-level resource. Each resource list page has a Tasks tab that lists only the tasks for that type of resource. For example, you can use the list page for storage systems to view only the tasks that are associated with storage systems.

About this task

Procedure

1. In the menu bar, select a type of top-level resource.
For example, to view the tasks that are related to storage systems, go to Storage and select the type of storage system you want to monitor.
Exception: Provisioning tasks that provision network-attached storage (NAS) file shares to servers or hypervisors are not shown on the resource list pages for servers or hypervisors. You must use the Tasks page to view these tasks.
2. Click the Tasks tab to view a list of all the tasks for the monitored resources of that type.
The information about tasks is organized into columns. These columns include information about the status of tasks, the most recent date and time when tasks ran, and the schedule for tasks.
Tip: The status icon on the Tasks tab represents the most critical status of the tasks that are associated with the monitored resources. For example, if one of the tasks for storage systems failed, the following icon is shown on the tab:
 Tasks
3. Optional: View the status icons on the Tasks page for a summary of task statuses.
This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.
Tip: To view descriptions of the icons and columns on the Tasks page, click the Help icon .


Viewing tasks for a specific resource

Use resource detail pages to view only the tasks for a specific monitored resource. Each resource detail page has a Tasks link that lists only the tasks for the associated storage system, server, hypervisor, switch, or fabric.

Procedure

1. In the menu bar, select a type of top-level resource.
For example, if you want to view the tasks for a storage system, go to Storage and select the type of storage system you want to monitor.
2. Right-click a resource in the list and select View Details.
The resource detail page is displayed. In the General section, the number next to Tasks represents the number of tasks that are associated with the resource. For example, Tasks (3) indicates that you can view three tasks.
The icon next to the Tasks link represents the most critical status of the tasks that are associated with the resource. For example, if a task failed, the following icon is shown:

Tasks (3) 

3. Click Tasks in the General section to view the tasks for the resource.
The information about tasks is organized into columns. These columns include information about the status of tasks, the most recent date and time when tasks ran, and the schedule for tasks.
4. Optional: View the status icons on the Tasks pane to view a summary of task statuses.
This summary includes the number of tasks that have a Failed, Warning, Successful, or Running status.
Tip: To view descriptions of the icons and columns on the Tasks pane, click the Help icon .


Viewing task details

Use task detail pages to view detailed information about the tasks that IBM Spectrum® Control uses to optimize resources and provision storage. You can use this information to troubleshoot the reason why a task failed, analyze details of a task before you implement task recommendations, or schedule future task runs.

About this task

On all the task detail pages, you can view information such as the task status, the duration and completion date of task runs, and the actions that you can take to manage the task. For tasks such as tiering analysis, balance analysis, and transform plan, you can view information about the recommendations that are generated by the most recent or earlier task runs. For provisioning tasks, you can view the status of the actions that the provisioning task completes for each of the resources that are configured in the task.

Procedure

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource.
For example, to view the tasks that are related to storage systems, go to Storage and select the type of storage system you want to monitor, and click the Tasks tab.
2. Locate the task that you want to view in more detail.
3. Right-click the task row and select View Details.
The task details page opens in a separate window. For example, the Tiering Analysis page opens if you select a tiering-analysis task row.
4. Optional: To view task logs, click Open Logs. You can use the information in the log file to troubleshoot any errors that might occur when a task is run.
5. Optional: For provisioning tasks, you can export information about the provisioning task to a text file for offline viewing. To export information, click the  icon at the top of the Provisioning page.

Viewing task logs




Task logs include detailed information about the status, actions, and progress of a task. Status information includes informational, warning, and error messages that are related to each action in a task that is taken during processing. You can use this information to troubleshoot any errors that might occur when a task is run.

About this task

The following examples show the type of information that you can view in the task logs:

- View the number of recommendations that are generated by tiering-analysis, balance-analysis, or transform-plan task runs.
- View the options that you have when you transform storage, analyze tiering, or balancing pools. For example, when you create a tiering-analysis task and the options that you choose are recorded in the task log.
- View error messages that are related to task processing.

Procedure

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource.
For example, to view the tasks that are related to storage systems, go to Storage and select the type of storage system you want to monitor, and click the Tasks tab.
2. Locate the task that has a log file that you want to view.
3. Right-click the task row and select Open Logs.
The Logs page shows log entries for the most recent task run.
4. Optional: On the Logs page, to view the log for a previous task run, select a task run from the Select a log list.
5. Optional: To view only the actions in a task run that have a Warning or Error status, select an option from the Show all list.
You can choose to view only actions that have the following statuses:
 -  Only error entries
 -  Only warning entries
 -  Error and warning entries
6. Optional: To view an explanation of the message that is associated with an action, click the link in the ID column.

Results

The information on the Logs page is automatically updated every 30 seconds. New entries are added to the end of a log. You can view the following information on the Logs page:

- The overall status of a task. The icon that is shown in the Select a log list represents the most critical status that was generated by an action in the task run.
- The status for each action in a task.
- The date and time when an action was completed. The date, time, and time zone of the action is shown in the Date and Time column.
- The ID of the message that is associated with an action. You can click the value in the ID column to view more information about a message.
- The description of an action.

Setting the number of task runs that are displayed

You can change the number of task runs that are shown in the GUI by configuring the settings on the History Retention page. By default, the information for the last five runs of a task is shown.

Procedure

1. In the menu bar, go to Settings > History Retention.
2. Click Edit.
3. In the Job logs field, enter the number of runs that you want to retain for display.
4. Click Save to save your changes.

Results

After you update the value for job logs, the list of task runs is automatically updated within a few minutes. For more information about retention settings, see [Configuring history and data retention](#).

Managing tasks for provisioning

Provisioning tasks are used to create storage volumes and assign the volumes to servers or hypervisors. Provisioning tasks are also used to create network-attached storage (NAS) file shares. Use the Tasks page in the GUI to manage all the tasks that are used by IBM Spectrum® Control to provision storage. Use resource list and resource details pages to manage the provisioning tasks for specific resources and resource types.

Before you begin

Discontinued support: Provisioning is no longer supported in IBM Spectrum Control. While the feature might still work in this release, it's recommended that you use another tool for your provisioning needs, when possible. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

About this task

Use the GUI to complete the following provisioning actions:

- Create provisioning tasks
- Run provisioning tasks immediately
- Schedule provisioning tasks to run at a specified date and time
- Export information about provisioning tasks to text files
- Delete provisioning tasks
- [Creating provisioning tasks](#)
You can create provisioning tasks with the GUI and assign storage volumes or NAS file shares to servers or hypervisors.
- [Running provisioning tasks](#)
Use the Tasks page to run provisioning tasks. You can run provisioning tasks immediately or you can schedule tasks to run at a date and time that you specify. When you run a provisioning task, the configuration changes that are contained in the task are implemented. For example, storage volumes are created and assigned to servers.
- [Exporting provisioning task information](#)
Use the Provisioning detail page to export details of a provisioning task to a text file. The task details that are shown on the Provisioning page are saved to the generated text file.

Related concepts

- [Provisioning storage overview](#)

Creating provisioning tasks

You can create provisioning tasks with the GUI and assign storage volumes or NAS file shares to servers or hypervisors.

Before you begin

You must have Administrator privileges to create provision tasks.

Related tasks

- [Provisioning storage with the IBM Spectrum Control GUI](#)
- [Collecting information about shares on storage systems](#)

Running provisioning tasks

Use the Tasks page to run provisioning tasks. You can run provisioning tasks immediately or you can schedule tasks to run at a date and time that you specify. When you run a provisioning task, the configuration changes that are contained in the task are implemented. For example, storage volumes are created and assigned to servers.

Before you begin

To run provisioning tasks, you must be the owner of the task or have Administrator privileges.

About this task

Provisioning tasks can run only once. You can run the task when you complete the steps for storage provisioning or you can save the provisioning task and run the task later.

The Run Now and Schedule actions are available for a provisioning task when the task is not running, and was not previously run.

Procedure

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource.
For example, to view the provisioning tasks that are related to servers, go to Servers_>Servers, and click the Tasks tab.
 2. Locate the provisioning task that you want to run.
 3. Depending on when you want the provisioning task to run, choose one of the following options:
 - To run the provisioning task immediately, right-click the task row and select Run Now. The task status changes to Running and the provisioning task runs immediately.
 - To schedule the provisioning task to run at a specific date and time, complete the following steps:
 - Right-click the task row and select Schedule.
 - From the Schedule page, select a future date and time from the Provision lists, and click Save.
- You can also run or schedule the provisioning task from the Provisioning detail page. On the Provisioning detail page, click Execute to run the task immediately or click Schedule to schedule the task to run at a specific date and time.

Results

The provisioning task runs immediately or on the scheduled date and time. To check the progress of the task, you can take the following actions:

- View the task status on the Tasks page.
- View detailed information about the progress of the task on the Provisioning detail page. You can view the task duration and the status of the actions that the provisioning task completes for each of the resources that are configured in the task.
- View the task logs from the Tasks page or from the Provisioning detail page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.


Exporting provisioning task information

Use the Provisioning detail page to export details of a provisioning task to a text file. The task details that are shown on the Provisioning page are saved to the generated text file.

About this task

If you do not have authority to provision storage, you can export the provisioning task details to a file and email the file to a user with Administrator privileges. The administrator can run or schedule the provisioning task if they approve the task details.

Procedure

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource.
For example, to view the provisioning tasks that are related to servers, go to Servers_>Servers, and click the Tasks tab.
2. Locate the provisioning task that you want to export.
3. Right-click the task row and select View Details.
4. From the Provisioning page, to start the export, click the  icon in the upper-right corner of the page.
5. Save the generated text file.
Depending on how your web browser is configured, you can specify the file name and the location where the file is saved.

Results

The provisioning task information is exported to a text file. The generated file includes all the provisioning task details that are listed on the Provisioning page. Columns that are hidden on the Provisioning page are included in the generated file. The information in the file is organized according to a default order and is not affected by the column order that is defined on the Provisioning page.

Managing tasks for tiering storage, balancing pools, and transforming storage

To optimize the resources in your storage environment, use tiering storage, balancing pools, and transforming-storage tasks. Use the Tasks page to manage all the tasks that are used to optimize storage. Use resource list and resource detail pages to manage the optimization tasks for resource types and for specific resources.

Before you begin

Discontinued support: Storage optimization is no longer supported in IBM Spectrum® Control. While the feature might still work in this release, it's recommended that you use another tool for your balancing and tiering needs, when possible. The storage transformation feature is still available. For a complete list of discontinued features, see [Discontinued features in IBM Spectrum Control](#).

To schedule, run, edit, or delete tiering storage, balancing pools, and transforming-storage tasks, you must have Administrator privileges.

About this task

To complete any of the optimization tasks, such as balancing pools, you must complete the following steps:

1. Create the analysis task with the GUI. For example, to balance pools create a balance-analysis task.
2. Run the analysis task to generate the optimization recommendations. For example, run the balance-analysis task to generate the recommendations to move volumes.
3. Edit analyze tiering and balancing pool tasks by modifying the thresholds.
4. Implement the recommendations to optimize storage. For example, implement the recommendations to balance pools by running the analysis-execution task.

Use the Tasks, resource list, and resource detail pages to take the following actions for optimization tasks:

- Run the optimization analysis immediately to generate optimization recommendations.
- Schedule the optimization analysis to run at a specified date and time.
- Implement optimization recommendations immediately.
- Schedule the implementation of optimization recommendations for a specified date and time.
- Pause, cancel, and resume the implementation of optimization recommendations.
- Delete optimization tasks.
- View task logs.
- [Creating analysis tasks](#)
To create tasks for optimizing resources in your storage environment, use IBM Spectrum Control GUI.
- [Running analysis tasks](#)
To run tiering analysis, balance analysis, and transform-plan tasks, use the Tasks page. When you run the analysis for any of these optimization tasks, recommendations are generated to optimize the storage resources that are configured in the task. For example, when you run a tiering-analysis task, recommendations are generated to tier volumes based on the criteria that you set in tiering policies.
- [Implementing recommendations to optimize storage](#)
Use the Tasks page and task detail pages to implement optimization recommendations to tier storage, balance pools, and transform storage. You can implement recommendations immediately or you can schedule the implementation of the recommendations for a future date and time.
- [Pausing, resuming, and canceling the implementation of recommendations](#)
IBM Spectrum Control uses analysis-execution tasks to implement optimization recommendations to tier storage, balance pools, and transform storage. To pause, resume, and cancel analysis-execution tasks, use the Tasks page.

Related tasks

- [Modifying the criteria for analyzing tiering](#)
- [Modifying the criteria for balancing pools](#)

Creating analysis tasks

To create tasks for optimizing resources in your storage environment, use IBM Spectrum® Control GUI.

Before you begin

You must have Administrator privileges to create tasks for optimizing resources in your storage environment.

About this task

Storage optimization tasks are created when you take any of the following actions in the GUI:

Tier storage

You can analyze tiering to move volumes to higher or lower tiers; a tiering-analysis task is created.

Balance pools

You can balance pools to distribute the workload of volumes across pools on the same tier; a balance-analysis task is created.

Transform storage

You can transform storage to move or convert volumes; a transform-plan task is created.

Related tasks

- [Tiering volumes by I/O density and I/O rate](#)
- [Transforming and migrating volumes](#)

Running analysis tasks

To run tiering analysis, balance analysis, and transform-plan tasks, use the Tasks page. When you run the analysis for any of these optimization tasks, recommendations are generated to optimize the storage resources that are configured in the task. For example, when you run a tiering-analysis task, recommendations are generated to tier volumes based on the criteria that you set in tiering policies.

Before you begin

To run tiering analysis, balance analysis, and transform-plan tasks, you must have Administrator privileges.

About this task

You can run tasks immediately or schedule tasks to run at a future date and time.

Tip: The Run Now action is available when the task is enabled and the task is not running.

Procedure

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource.
For example, to view the tiering-analysis tasks that are related to servers, go to Servers > Servers, and click the Tasks tab.
2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to run.
3. Right-click the task row and select Run Now.
The task status changes to Running.
4. Optional: To schedule future task runs for tiering analysis or balance-analysis tasks, right-click the task row and select Schedule Analysis. On the Schedule Analysis page, select the time that the task is run and how often the task is run.
Restriction: You cannot schedule future task runs for transform-plan tasks.
5. Optional: To enable or disable a tiering analysis or balance-analysis task, select Enabled or Disabled on the Schedule Analysis page.
When the task is enabled, the task runs according to the defined schedule. If the task is disabled, you can modify the task schedule but the task cannot run at the scheduled time.

Results

When the task run completes, you can view the recommendations that are generated on the task detail page. For example, you can view the recommendations that are generated for a tiering-analysis task on the Tiering Analysis page. To check the progress of the task run, you can take the following actions:

- View the task status on the Tasks page.
- View information about the progress of the task on the task detail page. You can view the duration of the task run, and the date and time that the task run completed.
- View the task logs from the Tasks page or from the task detail page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

Related tasks

- [Modifying the criteria for analyzing tiering](#)

Implementing recommendations to optimize storage

Use the Tasks page and task detail pages to implement optimization recommendations to tier storage, balance pools, and transform storage. You can implement recommendations immediately or you can schedule the implementation of the recommendations for a future date and time.

- **[Implementing recommendations immediately](#)**
To implement optimization recommendations to tier storage, balance pools, and transform storage, use the task detail pages. For example, use the Balance Analysis page to implement the recommendations to balance pools.
- **[Scheduling the implementation of recommendations](#)**
To schedule the implementation of optimization recommendations to tier storage, balance pools, and transform storage, use the Tasks page.

Related tasks

- [Modifying the criteria for analyzing tiering](#)

Implementing recommendations immediately

To implement optimization recommendations to tier storage, balance pools, and transform storage, use the task detail pages. For example, use the Balance Analysis page to implement the recommendations to balance pools.

Before you begin

To implement storage optimization recommendations, you must have Administrator privileges.

About this task

IBM Spectrum® Control uses analysis-execution tasks to implement optimization recommendations. When you implement recommendations for a tiering analysis, balance analysis, or transform-plan task, the following process takes place:

- An analysis-execution task is automatically created so that the recommendations can be implemented.
- The analysis-execution task runs immediately and the optimization recommendations are implemented. For example, volumes are moved to higher or lower tiers when storage-tiering recommendations are implemented.

Tip: In certain instances, running several optimization recommendations at the same time can cause one or more of the individual recommendation executions to encounter a Pool Exhaustion Exception. In these instances, the executions error out, rather than waiting for resources in the thread pool to become free. You might need to manually finish the execution of some recommendations, in particular, the old VDisk copy might not be removed, and you must remove it manually. To reduce the likelihood of encountering the issue again, contact IBM® Support for assistance in decreasing the total number of concurrent executions allowed.

You can view analysis-execution tasks on the Tasks page. The Related Task column shows the related tiering analysis, balance analysis, or transform-plan task.

The Execute action is available under the following conditions:

- Recommendations were generated by the latest run of the tiering analysis, balance analysis, or transform-plan task.
- An analysis-execution task was not already run for the latest set of recommendations.
- An analysis-execution task is not running or paused.

Procedure

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource.
For example, to view the analysis-execution tasks that are related to servers, go to Servers > Servers, and click the Tasks tab.
2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to implement.
3. Right-click the task row and select View Details.
For example, the Balance Analysis page opens if you select a balance-analysis task.
4. Click Execute.
An analysis-execution task is created and immediately implemented.

Results

The analysis-execution task runs immediately. To check the progress of the task, you can take the following actions:

- View the task status on the Tasks page.
- View the status of individual task recommendations on the Analysis Execution page.
- View the task logs from the Tasks page or from the Analysis Execution page. You can view detailed informational, warning, and error messages that are related to the task. You can use this information to troubleshoot any errors that might occur when a task is run.

Scheduling the implementation of recommendations

To schedule the implementation of optimization recommendations to tier storage, balance pools, and transform storage, use the Tasks page.

Before you begin

To schedule the implementation of storage optimization recommendations, you must have Administrator privileges.

The Schedule Execution action is available under the following conditions:

- Recommendations were generated by the latest run of the tiering analysis, balance analysis, or transform-plan task.
- An analysis-execution task is not running or paused.
- An analysis-execution task was not already run for the latest set of analysis recommendations. For each run of the tiering analysis, balance analysis, or transform-plan task, you can complete only one run of the analysis-execution task.

About this task

When you define the schedule, an analysis-execution task is created for the implementation of the optimization recommendations.

Tip: The analysis-execution task implements the recommendations that were current when the analysis-execution task was defined. For example, you define a tiering-analysis task that runs daily. You review the recommendations every day and, on a Tuesday, you schedule an analysis-execution task to run on the following Friday. When the analysis-execution task runs on Friday, the recommendations that were current on Tuesday are implemented, even if more recent recommendations are available.

Procedure

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource.
For example, to view the tiering analysis tasks that are related to servers, go to Servers > Servers, and click the Tasks tab.
2. Locate the tiering analysis, balance analysis, or transform-plan task that you want to schedule.
3. Right-click the task row and select Schedule Execution.
4. On the Schedule Execution page, select the date and time that the recommendations are implemented and click Save.

Results

An analysis-execution task is created that you can monitor on the Tasks page. The name of the analysis-execution task is based on the name of the related tiering analysis, balance analysis, or transform-plan task. The Related Task column shows the related tiering analysis, balance analysis, or transform-plan task.

When the analysis-execution task runs on the scheduled date, the optimization recommendations are implemented. For example, volumes are moved to higher or lower tiers when storage-tiering recommendations are implemented.

Tip: In certain instances, running several optimization recommendations at the same time can cause one or more of the individual recommendation executions to encounter a Pool Exhaustion Exception. In these instances, the executions error out, rather than waiting for resources in the thread pool to become free. You might need to manually finish the execution of some recommendations, in particular, the old VDisk copy might not be removed, and you must remove it manually. To reduce the likelihood of encountering the issue again, contact IBM® Support for assistance in decreasing the total number of concurrent executions allowed.

Pausing, resuming, and canceling the implementation of recommendations

IBM Spectrum® Control uses analysis-execution tasks to implement optimization recommendations to tier storage, balance pools, and transform storage. To pause, resume, and cancel analysis-execution tasks, use the Tasks page.

Before you begin

To pause, resume, or cancel analysis-execution tasks, you must have Administrator privileges.

About this task

Procedure

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource.
For example, to view the analysis-execution tasks that are related to servers, click Servers > Servers, and click the Tasks tab.
2. Locate the analysis-execution task that you want to pause, resume, or cancel.
3. Right-click the task row and select Pause, Resume, or Cancel.

The Pause action is only available when the analysis-execution task is running.

The Resume action is only available when the analysis-execution task is paused.

The Cancel action is available when the task is running or paused.

Results

When you pause, resume, or cancel analysis-execution tasks, the following processing occurs:

Pause

The status of the analysis-execution task changes to Paused and the recommendations that are being implemented are suspended.

Resume

The status of the analysis-execution task changes to Running and the recommendations are resumed.

Cancel

The status of the analysis-execution task changes to Canceled and the recommendations that are being implemented are stopped. Any recommendations that are already implemented for this analysis-execution task are not rolled back. For example, a recommendation is generated to move a volume from **pool_a** to **pool_b**. If this recommendation is already implemented before the task is canceled, the volume remains in **pool_b**. Use the Analysis Execution page to determine which recommendations were implemented for the task before the task was canceled.

You cannot restart an analysis-execution task that is canceled. If you want to implement the same set of recommendations again, you must complete the following steps:

1. Delete the analysis-execution task that you canceled.
2. Use the Schedule Execution action for the tiering analysis, balance analysis, or transform-plan task to create an analysis-execution task.

952

Renaming tasks

You can change the name of a task by using the task detail page. For example, use the Provisioning page to change the name of a provisioning task. A task name is automatically generated when the task is created. If the original name is ambiguous or is not clear enough, you can change the task name to a unique task name that you provide.

Before you begin

You must have the following privileges to rename the different task types:

Provisioning

To rename provisioning tasks, you must be the owner of the task or have Administrator privileges.

Other task types

To rename all other tasks, such as tiering analysis, balance analysis, transform plan, and analysis-execution tasks, you must have Administrator privileges.

Procedure

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource.
For example, to view the tasks that are related to servers, go to Servers > Servers, and click the Tasks tab.

2. Locate the task that you want to rename.
3. Right-click the task row and select View Details.
The task detail page opens in a separate window. For example, the Provisioning page opens if you select a provisioning task.
4. Type the new name in the task name field.
The naming conventions for task names are:
 - The maximum length is 64 characters.
 - Uppercase and lowercase letters are allowed.
 - Spaces and numbers are allowed.
 - The following special characters are allowed:
! # % & * + - / = ? ^ _ { } ()
5. Click Rename.

Results

The task is renamed. You can view the renamed task on the Tasks page.

Deleting tasks

Use the Tasks page to delete tasks, such as provisioning, tiering analysis, balance analysis, transform plan, and analysis-execution tasks.

Before you begin

You must have the following privileges to delete the different task types:

Provisioning

To delete provisioning tasks, you must be the owner of the task or have Administrator privileges.

Other task types

To delete all other tasks, such as tiering analysis, balance analysis, transform plan, and analysis-execution tasks, you must have Administrator privileges.

About this task

The following conditions apply for the different task types:

Analysis-execution

You can delete analysis-execution tasks that were canceled, already run, or scheduled to run in the future. You cannot delete tasks that are running or paused.

Other task types

You can delete all other tasks if the tasks were already run or if the tasks are scheduled to run in the future. For example, you can delete provisioning, tiering analysis, balance analysis, and transform-plan tasks that were already run. You cannot delete tasks that are running.

Procedure

1. From the menu bar, go to the Tasks page, a resource list page for a resource type, or a resource detail page for a specific resource.
For example, to view tasks that are related to servers, go to Servers, Servers, and click the Tasks tab.
2. Locate the task that you want to delete.
3. Right-click the task row and select Delete. Alternatively, you can delete the task by clicking Delete on the task detail page.
The Delete page opens in a separate window.
4. If you are deleting a provisioning task that was not yet run, enter the reason that you are deleting the task.
5. Click Delete.
To exit the window without deleting the task, click Cancel.

Planning copy data resources

To plan for your block storage requirements, you need to be able to see how you currently use block storage in your storage environment.

You can use the Copy Data view the following aspects of your block storage environment:

- You can see how much of your current allocated capacity is taken up by both the primary and backup data.
- You can see the ratio of allocated capacity that is taken up by primary data and plan for backup data according to your security policies.
- You can analyze storage device relationships to identify how redundancy is affecting available capacity.
- You can view how each type of storage you use is deployed in your environment.
- You can identify the volumes that are not backed up.
- **[Monitoring copy data resources](#)**
You can monitor the use and type of block storage resources in Copy Data.

Monitoring copy data resources

You can monitor the use and type of block storage resources in Copy Data.

You can get a complete breakdown of how each type of storage is used across the block storage environment.

Remote Relationships

Detailed information on remote relationships includes the type of relationship that exists between the source and target storage systems.



Synchronous relationship

Host writes are delayed until the source receives confirmation that the data is written to the target volume.



Asynchronous relationship

Host writes can continue while data is being written to the target volume.

Where the type of volume is known, one of the following icons appears next to the volume in question.



Standard volume

A volume that is not thin-provisioned, compressed, encrypted, or encrypt-able.



Thin provisioned volume

The volume is only allocated the required amount of server space as required.



Compressed VDisk volume

Data is compressed as it is written to the VDisk.

Consistency Group

Details the status and type of the consistency group, and the number of relationships that exist in the group.

FlashCopy

Details the information on the last FlashCopy® operation between each source and target storage system.

VDisk Mirrors

Details the information on the volumes, pools, tiers, and synchronization status of the VDisk mirrors.

- Identify target and source mirror disks at a glance.
- Easily see the disks that are compressed, thin, or normal.

Safeguarded Copy

For DS8000®, you can view the volumes that are backed up, the pool that the copies of a volume are written to, and the amount of capacity that is used to store the volume copies.

HyperSwap

Details the information on the status of the master and auxiliary volumes, pools, and tiers of the HyperSwap® relationships.

- Quickly identify the master and auxiliary volumes and pools.
- Drill down to see the amount of allocated disk for each volume.

Unprotected Volumes

Details the information on the volumes that are not backed up.

IBM Spectrum Control REST API

You can use the Representational State Transfer (REST) API for IBM Spectrum® Control to access information about resources and to generate custom capacity, configuration, and performance reports.

To get the information that you need about your resources, you can use a REST command line utility or you can use a web browser.

- [Retrieving data about resources by using a REST API command line utility](#)
You can connect to the REST API for IBM Spectrum Control and retrieve data by using a REST command line utility.
- [Retrieving data by using REST API with a web browser](#)
A convenient way to use the IBM Spectrum Control REST API is within the interface itself. You can quickly access resource information for your reporting capabilities by using a web browser.

Retrieving data about resources by using a REST API command line utility

You can connect to the REST API for IBM Spectrum® Control and retrieve data by using a REST command line utility.

You can use any REST command line utility, for example, GNU Wget. To access information about the GNU Wget utility, go to <https://www.gnu.org/software/wget/>. The IBM Spectrum Control REST API is hosted here: `https://<hostname>:9569/srm/REST/api/v1/`.

You can enter a command to authenticate with the REST API and store a security token in a file to use in your later queries, for example, `cookies.txt`.

You can use the utility, by entering the following command:

```
wget --post-data "j_username=<user name>&j_password=<password>"
--no-check-certificate --keep-session-cookies --save-cookies cookies.txt
https://<hostname>:9569/srm/j_security_check
```

You can use the security token in all subsequent commands that are issued against the REST API, by entering the following command:

```
wget --no-check-certificate --load-cookies cookies.txt
https://<hostname>:9569/srm/REST/api/v1/
```

Example (partial listing)

```
},
{
  "Description": "Provides a list of ports belonging to a parent storage system.
Format: \StorageSystems\<id>\Ports For a list of ports associated with a
port owner, specify the owner type and id in the URL.
```

To see a list of all storage systems, you can append the URL:

```
https://<hostname>:9569/srm/REST/api/v1/StorageSystems
```

To see more information about a resource type, you can further append the URL (where `<Name>` is the resource type, such as `StorageSystems`, `Switches`, `Servers`, etc.)

```
https://<hostname>:9569/srm/REST/api/v1/<Name>
```

To see more information about a specific resource, add the ID to the URL (where `<id>` is available from `https://<hostname>:9569/srm/REST/api/v1/<Name>`):

```
https://<hostname>:9569/srm/REST/api/v1/<Name>/<id>
```

To query for volumes that belong to a specific storage system, enter:

```
https://<hostname>:9569/srm/REST/api/v1/StorageSystems/57909/Volumes
```

Example (partial listing)

```
[
  {
    "Acknowledged": "No",
    "Allocated Space": "1.00",
    "Capacity": "1.00",
    "Controller": "Node 0",
    "Copy Relationship": "",
    "Easy Tier": "Tiered Pools\No",
    "Encryption": "",
    "Enterprise HDD Capacity": "1.00",
    "Format": "FB",
    "Hosts": "18",
    "LSS or LCU": "00",
    "Last Data Collection": "Apr 10, 2016, 15:03:57",
    "Name": "CET_RA_vol",
    "Nearline HDD Capacity": "",
    "Physical Allocation": "Fully Allocated",
    "Pool": "General Use 1",
    "RAID Level": "RAID 5",
    "SSD Capacity": "",
    "Service Class": "",
    "Shortfall": "",
    "Status": "Normal",
    "Storage System": "DS8000-2107-75BLG91-IBM",
    "Storage Virtualizer": "None",
    "Thin Provisioned": "No",
    "Ticket": "",
    "Tier Distribution": "0",
    "Unallocated Space": "0.00",
    "Unused Space": "",
    "Used Allocated Space": "",
    "Used Space": "1.00",
    "Virtualizer Disk": "None",
    "Volume ID": "0000",
    "Volume Number": "0",
    "Volume Unique ID": "75blg91\0000",
    "id": "79852"
  },

```

The information that is provided in the `Format` field in the example shows what combinations are possible. For example,

```
https://<hostname>:9569/srm/REST/api/v1/StorageSystems/<id>/Ports
```

Tips:

- If the `Format` field is omitted, you can add the ID of one resource to the URL.
- The name elements in the URL are case-sensitive.
- The names of resource types are plural.
- If the URL is not used correctly, you might receive the following error message:

```
{ "result": { "type": "E", "msgId": "BPCUI0099E", "time": "Apr 4, 2016 16:25:07", "text": "The storage resource is not available." }}
```

- If the URL returns a blank page, it means that no data is available for the resource. For example, if you query an application for file shares and there are no file shares in the application, the resulting page is blank.

Table 1. IBM Spectrum Control REST API URL examples

IBM Spectrum Control REST API	Description
<code>https://<hostname>:9569/srm/REST/api/v1/StorageSystems</code>	A list of all monitored storage systems.
<code>https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002</code>	Specific information for one storage system (ID is available from this URL: <code>https://<hostname>:9569/srm/REST/api/v1/StorageSystems</code>).
<code>https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes</code>	A list of all volumes for one storage system.
<code>https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance</code>	The available performance metrics for a volume.
<code>https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/RemoteReplication</code>	The remote mirror replication relationships for storage systems.
<code>https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance/824?granularity=sample&startTime=1455818400000&endTime=1455904800000</code>	Sample performance granularity information and a start and end time for the performance data.

Note: IBM® Storage is changing the capacity terminology that is used in IBM Storage products to make it more consistent. The new capacity terminology is used in IBM Spectrum Control 5.3.6 or later. The new terminology is used in the GUI, but the names used in the REST API have not changed. The output from the REST API continues to use Allocated Space instead of Used Capacity, Physical Allocation instead of Used Capacity %, and so on.

Retrieving data by using REST API with a web browser

A convenient way to use the IBM Spectrum® Control REST API is within the interface itself. You can quickly access resource information for your reporting capabilities by using a web browser.

To use the REST API with a web browser, use these steps:

1. Log on to the IBM Spectrum Control GUI. To start IBM Spectrum Control, see [Starting IBM Spectrum Control](#).
2. Open a new tab in your web browser.
3. Enter `https://<hostname>:9569/srm/REST/api/v1/`.

For example, to see a list of fabrics, enter `https://<hostname>:9569/srm/REST/api/v1/Fabrics` in a new tab of your web browser (partial listing).

```
[
{
  "Acknowledged": "No",
  "Active Zone Set": "",
  "Connected Switch Ports": "0",
  "Custom Tag 1": "",
  "Custom Tag 2": "",
  "Custom Tag 3": "",
  "Data Source Count": "1",
  "Fabric Type": "Cisco",
  "Last Successful Probe": "N/A",
  "Links": "0",
  "Location": "",
  "NPV Switches": "0",
  "Name": "VSAN0002",
  "Parent Fabric": "san_director",
  "Principal Switch of Fabric": "tpc-70swt-csc",
  "Probe Schedule": "N/A",
  "Probe Status": "N/A",
  "Status": "Error",
  "Switch Ports": "0",
  "Switches": "1",
  "Virtual": "Yes",
  "WWN": "2002000DECAC5081",
  "id": "2663"
}
```

To see the information for a specific fabric in your environment, append the URL with the ID:

`https://<hostname>:9569/srm/REST/api/v1/Fabrics/2017`

Example (partial listing)

```
{
  "Acknowledged": "No",
  "Active Zone Set": "",
  "Connected Switch Ports": "18",
  "Custom Tag 1": "",
  "Custom Tag 2": "",
  "Custom Tag 3": "",
  "Data Source Count": "2",
  "Fabric Type": "Cisco",
  "Last Successful Probe": "N/A",
  "Links": "4",
  "Location": "",
  "NPV Switches": "0",
  "Name": "san_director",
  "Parent Fabric": "",
  "Principal Switch of Fabric": "",

```



```

"Probe Schedule": "N\\A",
"Probe Status": "N\\A",
"Status": "Error",
"Switch Ports": "104",
"Switches": "2",
"Virtual": "No",
"WWN": "",
"id": "2017"

```

To see the switches that belong to the fabric "id": "2017", enter:

<https://<hostname>:9569/srm/REST/api/v1/Fabrics/2017/Switches>

Example (partial listing)

```

[
{
  "Acknowledged": "No",
  "Connected Fabrics": "",
  "Custom Tag 1": "",
  "Custom Tag 2": "",
  "Custom Tag 3": "",
  "Data Source Count": "1",
  "Domain ID": "10",
  "Fabric": "unstable BRCD_1-11",
  "Firmware": "v6.4.3g",
  "IP Address": "9.11.91.241",
  "Last Successful Monitor": "None",
  "Last Successful Probe": "None",
  "Links": "",
  "Location": "",
  "Mode": "Native",
  "Model": "Brocade 4100",
  "Name": "mdm-y76-swt",
  "Parent Switch": "",
  "Performance Monitor Interval (min)": "",
  "Performance Monitor Status": "Disabled",
  "Ports": "",
  "Principal Switch of Fabric": "cjswitch4",
  "Probe Status": "Never Probed",
  "Serial Number": "1070029",
  "Status": "Normal",
  "Vendor": "IBM",
  "Virtual": "No",
  "WWN": "100000051E347790",
  "id": "60989"
}
]

```

Note: IBM Spectrum Control can connect to Brocade switches and fabrics either directly or by using Brocade Network Advisor. For Brocade switches and fabrics that IBM Spectrum Control connects to directly, the Data Source Count value is 0.

Table 1. IBM Spectrum Control REST API URL examples

IBM Spectrum Control REST API	Description
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems">https://<hostname>:9569/srm/REST/api/v1/StorageSystems	A list of all monitored storage systems.
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002	Specific information for one storage system (ID is available from this URL: <a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems">https://<hostname>:9569/srm/REST/api/v1/StorageSystems).
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes	A list of all volumes for one storage system.
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance	The available performance metrics for a volume.
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/RemoteReplication">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/RemoteReplication	The remote mirror replication relationships for storage systems.
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance/824?granularity=sample&startTime=1455818400000&endTime=1455904800000">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance/824?granularity=sample&startTime=1455818400000&endTime=1455904800000	Sample performance granularity information and a start and end time for the performance data.

Note: IBM® Storage is changing the capacity terminology that is used in IBM Storage products to make it more consistent. The new capacity terminology is used in IBM Spectrum Control 5.3.6 or later. The new terminology is used in the GUI, but the names used in the REST API have not changed. The output from the REST API continues to use Allocated Space instead of Used Capacity, Physical Allocation instead of Used Capacity %, and so on.

Reporting

Keep your colleagues and management up-to-date by sending inventory, capacity, performance, and storage consumption reports by email or by exporting information about your storage environment.

Reports overview

Share the information that is collected about your storage environment by sending reports by email or by exporting the information to create your own reports.

In IBM Spectrum® Control, you can share information with your colleagues and management:

- [By creating predefined inventory and capacity reports that you can save to your file system and send by email](#)
- [By creating custom reports that you can save to your file system or send by email](#)
- [By creating chargeback and consumer reports that you can send by email](#)
- [By using the Representational State Transfer \(REST\) API to get information about storage resources and to generate configuration and performance reports](#)
- [By exporting information about the configuration, capacity, and performance of storage resources to HTML, PDF, or CSV files](#)

More about exporting information to files:
You can use the export feature as follows:

- To export capacity and performance information about your storage resources from IBM Spectrum Control into spreadsheets and the reports that you create for your organization
- To share information with your colleagues about the storage resources in your environment
- To make your colleagues aware of issues and to help them investigate and resolve capacity and performance issues

You can use an external reporting tool, Cognos® Analytics reporting interface, to view predefined reports and create custom reports about IBM Spectrum Control. For more information, see [Reporting with Cognos Analytics](#).

Predefined, custom, and storage consumption reports

Reports



Predefined inventory and capacity reports

Leverage and share the insights that IBM Spectrum Control gains from the metadata that it collects and analyzes to generate information about your storage assets and to monitor capacity usage and plan capacity purchases.

➔ [Create predefined capacity reports](#)

➔ [Create predefined inventory reports](#)

Custom reports

From any table view for a resource, leverage and share the insights that IBM Spectrum Controls gains into the capacity, configuration, performance, and health status of the storage resources that you monitor.

➔ [Create custom reports](#)

Chargeback and consumer reports

Create, schedule, and send chargeback and consumer reports by email to make your organization aware of the cost and the amount of capacity that is used by storage consumers such as applications.

➔ [More information about chargeback and consumer reports](#)

➔ [Create chargeback reports](#)

➔ [Create consumer reports](#)

In the IBM Spectrum Control you can create the following types of reports, which you can configure, schedule, and send by email:

Predefined reports

You can use predefined reports:

- To notify colleagues about the capacity issues that affect them, and to monitor the capacity usage of resources, such as storage systems, pools, and hosts, to ensure that you have sufficient capacity to meet current usage and to help plan capacity purchases
- To generate inventory reports about your storage assets, such as storage systems, hypervisors, servers, IBM Spectrum Virtualize nodes, and IBM Spectrum Scale nodes

Custom reports

You use custom reports to create a variety of reports such as reports that provide information about your storage assets, and capacity and performance reports that provide information about the storage usage and performance of your resources.

Tip: You can open the table view for any resource, such as storage systems, volumes, pools, general groups, applications, and departments and click Create Report to create a custom report that you can save or send in a few simple steps. To create performance reports, open the Performance page for the resource and click Create Report.

Storage consumption reports

Chargeback reports make the owners of resources, such as applications, departments, hypervisors, or physical servers, aware of the amount and cost of the block and file storage that is consumed.

Consumer reports make the owner of a specific resource, such as an application, department, hypervisor, or physical server, aware of the amount and cost of the block and file storage that is consumed.

How to create reports

Watch a short video about how to create and customize a capacity report, attach a file to the report, and schedule when to run it in IBM® Storage Insights Pro.



- [Creating predefined capacity reports](#)
To share information about capacity availability, usage, and growth, create capacity reports. You can configure and refine the information that is included in the report so that your colleagues get the information that they need to monitor the resources that they manage and make capacity planning decisions.
- [Creating predefined inventory reports](#)
To share information with your colleagues about the configuration and properties of your storage assets, create inventory reports.
- [Creating custom reports](#)
Create, configure, and save or send custom reports by email that include asset, capacity, configuration, health status, or performance information about your storage resources.
- [Creating chargeback and consumer reports](#)
To help plan capacity purchases and make your organization aware of the cost and the amount of the storage that is used by storage consumers, create chargeback and consumer reports.
- [Running reports](#)
You can run reports that you created without defining a schedule or changing the schedule that was created for the report.
- [Editing reports](#)
Change the configuration and scheduling for reports.
- [Deleting reports](#)
Remove the reports that you don't need.
- [Configuring the email server](#)
To send reports, you must configure the email server.
- [Investigating issues with reports](#)
You can access log and trace files to investigate the issues that you might have when reports are generated and sent by email or saved to the file system.
- [Types of predefined capacity and inventory reports](#)
The predefined capacity and inventory reports that you can create for your storage resources are listed.
- [Reports FAQ](#)
Find answers to questions about the reports that you can create in IBM Spectrum Control.
- [Capacity metrics for reports](#)
To gain insights into storage usage, review the capacity metrics that are collected, analyzed, and shown in reports.
- [Using the REST API to generate reports](#)
Use the REST command line utility or a web browser to generate capacity, configuration, and performance reports.
- [Exporting information about resources](#)
Notify your colleagues about the current state and potential issues with your storage environment by exporting the capacity and performance information that is shown in the GUI.
- [Reporting with Cognos Analytics](#)
Use the Cognos Analytics reporting tool to view predefined reports and create custom reports about the resources that are monitored by IBM Spectrum Control.

Creating predefined capacity reports

To share information about capacity availability, usage, and growth, create capacity reports. You can configure and refine the information that is included in the report so that your colleagues get the information that they need to monitor the resources that they manage and make capacity planning decisions.

Before you start

If you want to send reports by email, you must set up the mail server. If you already set up the mail server to send alerts or reports, or if you want to only save reports to your file system, you don't have to set up the mail server.

If you want the reports that you will create to be saved to your file system to be stored in a folder other than the default reports folder, you must create the folder before you create the reports.

Saving reports to file systems on remote servers: The machine on which the Web server component was installed and runs must be able to access the remote or shared file system that is specified for saving reports. To make file systems on a server available to remote clients, they must be mounted as a CIFS export on Windows operating systems or as an NFS export on UNIX or Linux® operating systems. On Windows operating systems, the file system must be mounted as the local system user or from a service that is running under the local system account so that it is visible to all logon sessions. One way of mapping a drive as the local system user is to map the drive from a psexec command window. The psexec command window is available as part of the Sysinternals command-line tools for Windows operating systems.

Step 1: Pick your predefined report

Click Reports, and then click Create Report and pick the capacity report that you want to create.




Step 2: Name the report

Provide a unique name for your report. You can use alphabetical and numerical characters, hyphens, dashes, and blank spaces.

Step 3: Specify the scope of the report

Choose one of these options:

- Option 1: Generate capacity information about all of the resources of a specific type.
- Option 2: Generate capacity information about a selection of resources of a specific type.
- Option 3: Generate capacity information about the resources of a specific type in a group.

Option 1: All	Option 2: Selection	Option 3: In a Group
		
Include information in the report about all of the storage systems.	Include information in the report about a selection of the storage systems.	Include information in the report about the storage systems that are assigned to a general group, such as the production platform.

Option 1: All of the resources

Choose this option to generate a report about all of the resources, such as all storage systems, pools, volumes, managed disks, or NAS file systems.

Option 2: One or more of the resources

Choose this option to generate capacity information about a selection of resources such as a selection of storage systems, pools, volumes, managed disks, or NAS file systems.

When you create capacity reports about managed disks, pools, server , or volumes, you can group the types of resources, such as managed disks by storage systems, pools by storage systems, server by applications, and volumes by server .

Try it out: Click Reports, click Create Report, and choose the Pools capacity report. Choose capacity information for one or more resources, and select pools by storage systems or pools by tiers. To refine the list of the resources that are shown, you can enter a matching pattern for the names of the storage systems, or you can select the storage systems that contain the pools that you want to include in the report.


Predefined Capacity Reports: Selecting the resources that you're interested in

Pools capacity reports

Pools by storage systems



Pools by tiers




Volumes capacity reports

Volumes by storage systems



Volumes by pools



Volumes by hosts




Volumes by applications




Managed disks capacity reports

Managed disks by storage systems



Hosts capacity reports

Hosts by applications



Option 3: Resources assigned to groups

You can also generate capacity information about the storage systems that belong to a group that you created, such as a general group that consists of the storage systems that are assigned to your production platform, or the storage resources that are assigned to an application or department.

Tip: The report you pick, determines the level of information that is shown about applications in reports. For example, if you want to see the capacity information about the set of volumes that you added to an application, create a Volumes capacity report and select the application. Or, if you want to see the capacity information about the filesets that you added to an application, create a Filesets capacity report and select the application.

Step 4: Select columns

To ensure that the report's recipients get the information that they need, you can add more information, such as the location of the resources or custom tags. You can also remove the information that your colleagues don't need.

Step 5: Add filters

To refine the information in the report, you can add up to four filters.

Sample filters for block storage systems capacity reports

Sample 1: Filter by storage system name	Sample 2: Filter by Used Capacity %
Column Storage System	Column Used Capacity %
Comparison Contains	Comparison > =
Value storwize	Value 80
Only include capacity information about IBM Storwize family storage systems.	Only include capacity information for IBM Storwize family storage systems with used capacity values of 80% or higher.

For example, if you want to notify resource administrators or owners about capacity shortfalls, you can set threshold values for capacity information. Or if you want to exclude information that is not of interest to the report's recipients, you can add a filter.

Don't forget to check the table when you apply filters to make sure that the recipients of the report will get the information that they need.

Step 6: Schedule and deliver the report

To notify your colleagues of capacity or configuration anomalies, create one-off capacity and inventory reports or create a schedule to provide regular updates about the capacity or state of your storage assets.

Tip: Instead of sending the report by email, you can save the report to the default reports folder or to a folder on your file system. Alternatively, when you send the report by email, you can also attach the report as a CSV, PDF, or HTML file.

Review, edit, and delete reports

To review the report that you created, click Reports, expand the relevant section, such as Capacity Report for predefined capacity reports and Inventory Report for predefined inventory reports, and then select the report.

You can edit the report, such as change the name and scope of the report or schedule of the report. To delete the report, click Actions,⌵Delete Report.

Tip: You want to save a scheduled report to your file system or send it by email report now, but you don't want to change the original schedule? Click Reports. Select the report, and click Actions,⌵Run Now. The report is sent without changing the schedule that you saved.

- [Tutorial: Creating a predefined capacity report about storage systems](#)
Create, schedule, and send predefined capacity reports about your storage systems by email, to keep your colleagues informed about capacity usage and to help them plan capacity and prevent capacity shortages.
- [Tutorial: Creating a predefined capacity report about pools](#)
Create, schedule, and send predefined capacity reports about your pools by email, to keep your colleagues informed about capacity usage and to help them plan capacity purchases and prevent capacity shortages.
- [Tutorial: Creating a predefined capacity report about tiered pools](#)
Create a pools by tier report to share information about the capacity of the pools that are assigned to your production platform.
- [Tutorial: Creating a predefined capacity report about the volumes assigned to servers](#)
Create, schedule, and send predefined reports about the capacity of the volumes that are assigned to servers .
- [Tutorial: Creating a predefined capacity report about managed disks by storage systems](#)
Share information with your colleagues about the capacity of your managed disks grouped by the storage systems that they belong to.
- [Adding resources to applications to generate large reports](#)
To generate capacity reports about a large number of resources, add the resources to an application and then generate the report.

Tutorial: Creating a predefined capacity report about storage systems

Create, schedule, and send predefined capacity reports about your storage systems by email, to keep your colleagues informed about capacity usage and to help them plan capacity and prevent capacity shortages.




About this task

You can create a predefined report about the capacity usage of all of the storage systems in your storage environment, a selection of storage systems in your storage environment, or the storage systems that are assigned to a group such as a general group.

Procedure

1. Click Reports, and then click Create Report.
2. In the Capacity Reports section, click Block Storage Systems.
3. Enter the unique name of the report.
4. To specify the range of the report, choose one of the following options:

Specify the range of the report, choose one of the following options:

<div>Option 1: All</div> <div></div> <div>Include information in the report about all of the storage systems.</div>	<div>Option 2: Selection</div> <div></div> <div>Include information in the report about a selection of the storage systems.</div>	<div>Option 3: In a Group</div> <div></div> <div>Include information in the report about the storage systems that are assigned to a general group, such as the production platform.</div>
--	---	--

Option	Description
Capacity information about all of the storage systems	Generate capacity information about all of the block, file, and object storage systems in your storage environment.
Capacity information about one or more storage systems	Generate information about a set of storage systems.
Capacity information about the storage systems that are assigned to a group	Generate information about the storage systems that were assigned to a group such as the production, preproduction, or test platforms general group.

5. Add and remove columns to choose the capacity and any other information that you want to include in the report.
6. Add up to four filters to deliver the capacity information that the report's recipients need.

Sample 1: Filter by storage system name

Column
Storage System

Comparison
Contains

Value
storwize

Only include capacity information about IBM Storwize family storage systems.

Sample 2: Filter by Used Capacity %

Column
Used Capacity %

Comparison
> =

Value
80

Only include capacity information for IBM Storwize family storage systems with used capacity values of 80% or higher.

Tip: To make sure that the report's recipients get the information that you want them to get, use the report preview feature. Each time you apply a filter, or make a configuration change to the report, such as add or remove a column or filter, the table view of the report is refreshed.

7. Choose one of the following options:

Option	Description
Email	Save and send the email now or attach the report to the email as a CSV, PDF, or HTML file.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV, PDF, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

Results

The Reports page opens where you can preview and edit the report that you created.

Tutorial: Creating a predefined capacity report about pools

Create, schedule, and send predefined capacity reports about your pools by email, to keep your colleagues informed about capacity usage and to help them plan capacity purchases and prevent capacity shortages.

About this task

You want to send a regular report to the owners of block storage resources so that they'll know:

- Which pools are depleted
- When the pools will run out of capacity

To do this, you include the Zero Capacity column and create a filter.

What's zero capacity: The capacity information that is collected over 180 days is analyzed to determine, based on historical storage consumption, when the pools will run out of capacity. The pools that have already run out of capacity are marked as depleted. For the other pools, a date is provided so that you know when the pools are projected to run out of capacity. If sufficient information isn't collected to analyze the storage usage of the pool, *None* is shown as the value for zero capacity. If a capacity limit is set for the pool, the date shown in the Zero Capacity column is the date when the available capacity based on the capacity limit will be depleted. For example, if the capacity limit for a 100 GiB pool is 80%, it is the date when the available capacity of the pool is less than 20 GiB. *Depleted* is shown in the column when the capacity limit is reached.

Procedure

1. Click Reports, and then click Create Report.
2. In the Capacity Reports section, click Pools.
3. Enter the unique name of the report.
4. To specify the range of the report, choose one of the following options:

Option 1: All

Include information in the report about all of the pools in your storage environment.

Option 2: Selection

Include information in the report about a selection of pools.


Option 3: In a Group

Include information in the report about the pools that are assigned to a group, such as an application.

Option	Description
Capacity information about all of the pools	Generate capacity information about all of the pools in your storage environment.
Capacity information about one or more pools	Generate information about a set of pools. You can select the pools by the storage systems that they belong to or by the tiers that the pools are assigned to.
Capacity information about the pools that are assigned to a group	Generate information about the pools that were assigned to a general group such as an application.

5. Add and remove columns to choose the capacity and any other information that you want to include in the report.
6. Add up to four filters to deliver the capacity information that the report's recipients need.

Sample: Filter by Zero Capacity



Only include capacity information about the pools that are depleted or that could be analyzed to determine when the pools are projected to run out of capacity.

Tip: To make sure that the report's recipients get the information that you want them to get, use the report preview feature. Each time you apply a filter, or make a configuration change to the report, such as add or remove a column or filter, the table view of the report is refreshed.

7. Choose one of the following scheduling options:

Option	Description
Email	Save and send the email now. You can also attach the report to the email as a CSV, or HTML file, or both.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV file, PDF file, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

Results


The Reports page opens where you can preview and edit the report that you created.

Tutorial: Creating a predefined capacity report about tiered pools

Create a pools by tier report to share information about the capacity of the pools that are assigned to your production platform.

Before you begin

The pools that are used on your production platform are assigned to one or more tiers.

Assign pools to tiers: Click **Storage > Pools**. To add the Tier column to the table, right-click any column heading, and select **Tier**. Select and then right-click the pools on your production platform, and click **Set Tier**. If you used a naming pattern for the pools or storage systems that are assigned to your production platform, you can filter the pools by pool or storage system name. Just click . Then select **Name** or **Storage System** and enter the pattern that matches the names of the pools or storage systems in the filter.

About this task

You want to create a scheduled report that shows the key capacity values of the pools on your production platform.

Procedure

1. Click **Reports**, and then click **Create Report**.
 2. Pick the **Pools capacity** report.
 3. Enter the name of the report, and click **Capacity information** about one or more pools.
 4. Select **Tiers** from the **Pools** drop-down list.
A list of the tiers that were created is shown.
 5. Select the tiers that you created for your production platform.
When you click **Next**, capacity information about the pools assigned to the tiers that you selected is shown.
 6. Select the capacity information that you want to include in the report.
Tip: Include the name of the tier in the report. In the **General** section, click **Tier**. Drag the **Tier** column and drop it before the **Name** column. To sort the pools by the tiers they are assigned to, click the **Tier** column.
 7. To refine the information that you want to include in the report, add filters.
 8. Schedule the delivery of the report.
- [Optimizing storage tiering](#)
 - [Setting the tier level of storage pools](#)

Tutorial: Creating a predefined capacity report about the volumes assigned to servers

Create, schedule, and send predefined reports about the capacity of the volumes that are assigned to servers .

About this task


When you create reports about the capacity of volumes, you can select volumes:

- By the storage systems that they belong to.
- By the pools that they belong to.

- By the servers that they are assigned to.
- By the applications that they are assigned to.

In this tutorial, you create a capacity report about the volumes that are assigned to the servers in your storage environment.


Procedure

1. Click Reports, and then click Create Report.
2. Click the Volumes capacity report.
3. Enter the name of the report.
4. Select Capacity information about one or more volumes.
5. On the Select Volumes page, select Servers from the Volumes list.
6. On the Select Volumes by Servers page, you can:
 - a. Choose all of the servers.
 - b. Choose one or more of the servers.
 - c. Select Name from the Filter  list. Then, enter a pattern that matches the names of the servers that you want to select.
7. Choose the columns that you want to include in the report.
You can rearrange the columns by dropping the columns where you want them to appear in the report.
8. Add a filter and then click Apply Filter.
For example, you can create a weekly report that lists the volumes with a used capacity greater than 80%.
9. Schedule the report.

Tutorial: Creating a predefined capacity report about managed disks by storage systems

Share information with your colleagues about the capacity of your managed disks grouped by the storage systems that they belong to.

Procedure

1. Click Reports, and then click Create Report.
2. Pick the Managed Disks capacity report.
3. Enter the name of the report and choose Capacity information about one or more managed disks.
4. On the Select Managed Disks page, change the selection from Managed Disks to Block Storage Systems.
5. On the Select Managed Disks by Block Storage Systems page, you can:
 - a. Choose all of the storage systems that have managed disks.
 - b. Choose one or more of the storage systems that have managed disks.
 - c. Select Name from the  the column filter. Then, enter a pattern that matches the names of the storage systems that you want to select.
6. Choose the columns that you want to include in the report.
You can rearrange the order of the columns by dropping the columns where you want them to appear in the report.
7. Add and apply filters.
8. Schedule the report.

Adding resources to applications to generate large reports

To generate capacity reports about a large number of resources, add the resources to an application and then generate the report.

About this task

For example, you want to generate a capacity report about all of the file sets, hosts, or volumes in your production environment. To do this, you can add the resources to an application and then generate a report about the application.

You can also add resources to general groups to generate large reports. However, when you add or remove resources, you'll have to add or remove the resources from the general group that you created. Whereas, when you use naming patterns to match the resources that you associate with applications, the applications are automatically updated when resources are added or removed.

Procedure

1. To add file sets or volumes to an application, complete these steps:

Option	Description
File sets and volumes	<ol style="list-style-type: none"> a. Click Groups > Applications. b. Click Create Application, enter the name of the application, and then click Create. c. Click Assign Storage Resources to the Application. d. Choose a selection of volumes or file sets that match a pattern. e. To add all the file sets or volumes to the filter, type an asterisk (*) . f. Save the filter.

Option	Description
Hosts	<ol style="list-style-type: none"> Click Groups > Applications. Click Create Application, enter the name of the application, and then click Create. Click Assign Storage Resources to the Application. Choose all storage that belongs to set of servers or hypervisors. Enter a pattern that matches the names of you servers. Save the filter.

- Click Reports, and then click Create Report.
- Pick the capacity report that you want to create.
- Enter the name of the report, choose the option for groups and select the group that you created.
- Complete the steps for creating the report.

Creating predefined inventory reports

To share information with your colleagues about the configuration and properties of your storage assets, create inventory reports.

Before you start

If you want to send reports by email, you must set up the mail server. If you already set up the mail server to send alerts or reports, or if you want to only save reports to your file system, you don't have to set up the mail server.

If you want the reports that you will create to be saved to your file system to be stored in a folder other than the default reports folder, you must create the folder before you create the reports.

Saving reports to file systems on remote servers: The machine on which the Web server component was installed and runs must be able to access the remote or shared file system that is specified for saving reports. To make file systems on a server available to remote clients, they must be mounted as a CIFS export on Windows operating systems or as an NFS export on UNIX or Linux® operating systems. On Windows operating systems, the file system must be mounted as the local system user or from a service that is running under the local system account so that it is visible to all logon sessions. One way of mapping a drive as the local system user is to map the drive from a psexec command window. The psexec command window is available as part of the Sysinternals command-line tools for Windows operating systems.

Step 1: Pick your predefined report

Click Reports, and then click Create Report and pick the inventory report that you want to create.

Step 2: Name the report

Provide a unique name for your report. You can use alphabetical and numerical characters, hyphens, dashes, and blank spaces.

Step 3: Specify the scope of the report

Generate information about all of the resources of a specific type, such as an inventory report for block storage systems, or IBM Spectrum Virtualize nodes, or IBM Spectrum Scale nodes.






Step 4: Select columns

To ensure that the report's recipients get the information that they need, you can add more information, such as the location of the resources or custom tags. You can also remove the information that your colleagues don't need.

Step 5: Add filters

To refine the information in the report, you can add up to four filters.

Sample filters for block storage systems inventory reports

Sample 1: Filter by type of storage system			Sample 2: Filter by Firmware version		
					
Column Type		Comparison Contains	Column Firmware	Comparison Not Contains	Value 7.8.1.5
Only include inventory information for IBM Storwize family storage systems.			Only include information for IBM Storwize family storage systems with firmware versions that don't match the value that you enter.		

For example, if you want to check whether your IBM® Storwize® family block storage systems have the correct firmware version, you can add filters to list the storage systems with firmware versions that don't match the value that you enter.

Don't forget to check the table when you apply filters to make sure that the recipients of the report will get the information that they need.

Step 6: Schedule and deliver the report

To notify your colleagues of inventory issues, create one-off inventory reports or create a schedule to provide regular updates about the state of your storage environment.

Tip: Instead of sending the report by email, you can save the report to the default reports folder or to a folder on your file system. Alternatively, when you send the report by email, you can also attach the report as a CSV, PDF, or HTML file.

Review, edit, and delete reports

To review the report that you created, click Reports, expand the relevant section, such as Inventory Report, and then select the report.

You can edit the report, such as change the name and scope of the report or schedule of the report. To delete the report, click Actions...Delete Report.

Tip: You want to save a scheduled report to your file system or send it by email report now, but you don't want to change the original schedule? Click Reports. Select the report, and click Actions...Run Now. The report is sent without changing the schedule that you saved.

- **Tutorial: Creating an inventory report about block storage systems**

Create a predefined report that includes information such as the name, type, model, vendor, and location of your block storage systems.

Tutorial: Creating an inventory report about block storage systems

Create a predefined report that includes information such as the name, type, model, vendor, and location of your block storage systems.

Procedure

1. Click Reports, and then click Create Report.
2. In the Inventory Reports pane, choose Block Storage Systems.
3. Enter the unique name of the report.
4. Choose the inventory information that you want to include in the report.
5. Add up to four filters.

Tip: Don't forget to click Apply Filter when you add each filter, and check the report preview to ensure that the report's recipients will get the information that they need.

6. Choose one of the following scheduling options:

Option	Description
Email	Save and send the email now. You can also attach the report to the email as a CSV, PDF, or HTML file.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV, PDF, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

Results

The Reports page opens where you can preview and edit the report that you created.

Creating custom reports

Create, configure, and save or send custom reports by email that include asset, capacity, configuration, health status, or performance information about your storage resources.

Before you begin

If you want to send custom reports by email, you must set up the mail server. If you already set up the mail server to send alerts, or chargeback and consumer reports, or if you want to only save reports to your file system, you don't have to set up the mail server.

If you want the reports that you will create to be saved to your file system to be stored in a folder other than the default reports folder, you must create the folder before you create the reports.

Saving custom reports to file systems on remote servers: The machine on which the Web server component was installed and runs must be able to access the remote or shared file system that is specified for saving custom reports. To make file systems on a server available to remote clients, they must be mounted as a CIFS export on Windows operating systems or as an NFS export on UNIX or Linux® operating systems. On Windows operating systems, the file system must be mounted as the local system user or from a service that is running under the local system account so that it is visible to all logon sessions. One way of mapping a drive as the local system user is to map the drive from a psexec command window. The psexec command window is available as part of the Sysinternals command-line tools for Windows operating systems.

About this task

Use custom reports to quickly create reports to alert members of your organization about performance anomalies, health issues, capacity shortfalls, or capacity depletion. You can also schedule custom reports to keep a close watch on critical resources such as applications, or the storage resources that you use in your production environment.

From any table view that shows information about your storage systems or their related resources, you can create, configure, schedule, and save or send reports by email.

For example, you can create custom reports for top-level storage resources such as:

- Storage systems
- Hypervisors
- Servers
- Fabrics
- Switches

- Back-end storage systems

You can create custom reports for internal resources such as:

- Disks
- Pools
- Volumes
- Filesets
- File systems
- Network Shared Disks

And, if you added groups, you can create custom reports for general groups, applications, and departments.

File names for reports: The format for the name of reports that are saved to the file system is: <report name>.<epoch time>.<HTML|PDF|CSV>, where epoch time reflects the time of the server where IBM Spectrum® Control is installed. For example, the name of your report is *MyReport*, and you save it as a CSV file on 16 October, 2018 at 08:15pm (UTC) server time. The name of the file for the report is *MyReport_1539720900000.CSV*.



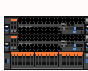

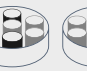
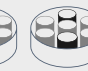
Tip: Monitor the default reports folder or the custom folders that you use to save reports to ensure that you have enough space for saving reports. For example, if you schedule and save a daily report to your file system, each version of the report is saved to the folder you choose. The factors that determine space usage are:

- The number of rows in the report.
- The frequency of the report.
- The format of the report.

Try it out: Open any page that shows information about your storage resources, such as the Block Storage Systems page, click Create Report, and create a custom report. Or, go to the Performance page and click Create Report, to create a performance report.

Procedure

1. Open any page that shows capacity or performance information about your storage systems, and then click Create Report.
2. Provide a unique name for your report. You can use alphabetical and numerical characters, hyphens, dashes, and blank spaces.
Remember: When you create your report, you can also use the table view features to customize your report, such as:
 - Drag the column headings to reorder the information that you want to show in the report.
 - Click the column heading in the table view to sort the values in the column.
3. To ensure that the report's recipients get the information that they need, you can add other information, such as the location of the resources or custom tags and remove the information they don't need.
4. To refine the capacity or asset information that you want to share, you can add up to four filters.

Sample 1: Filter by storage system name			Sample 2: Filter by Used Capacity %		
					
Column Storage System		Comparison Contains	Column Used Capacity %	Comparison > =	Value 80
Only include capacity information about IBM Storwize family storage systems.			Only include capacity information for IBM Storwize family storage systems with used capacity values of 80% or higher.		

For example, if you want to notify resource administrators or owners about capacity shortfalls, you can set threshold values for capacity information. Or if you want to exclude information that is not of interest to the report's recipients, you can add a filter. Check the report preview when you apply filters to make sure that the recipients of the report will get the information that they need.

Tip: You can add advanced filters to refine information about capacity and information about the assets in your storage environment. If advanced filters aren't available for a resource, you can use the filter feature that is available in the table. For example, in reports about volumes, you can filter by volume, pool, or storage system name, or by Easy Tier® attributes.

5. Choose one of the following scheduling options:

Option	Description
Email	Save and send the email now. You can also attach the report to the email as a CSV, PDF, or HTML file.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV, PDF, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

Results

To review the report that you created, click Reports, Reports, expand Customs Report, and then select the report.

You can edit the report, such as change the name and the schedule for the report. To delete the report, click Actions, Delete Report.

Tip: You want to save a scheduled report to your file system or send it by email report now, but you don't want to change the original schedule? Click Reports, Reports. Select the report, and click Actions, Run Now. The report is sent without changing the original schedule.

- [Tutorials: Creating custom capacity and performance reports for applications](#)
In these tutorials, you create an application so that you can create, schedule, and send reports by email about the storage resources that the application consumes.
- [Custom capacity and performance view reports](#)
When you create custom capacity or performance view reports, you can specify a relative time range for the information that is shown about the resources in the reports.

Related concepts

- [Monitoring the capacity of resources](#)

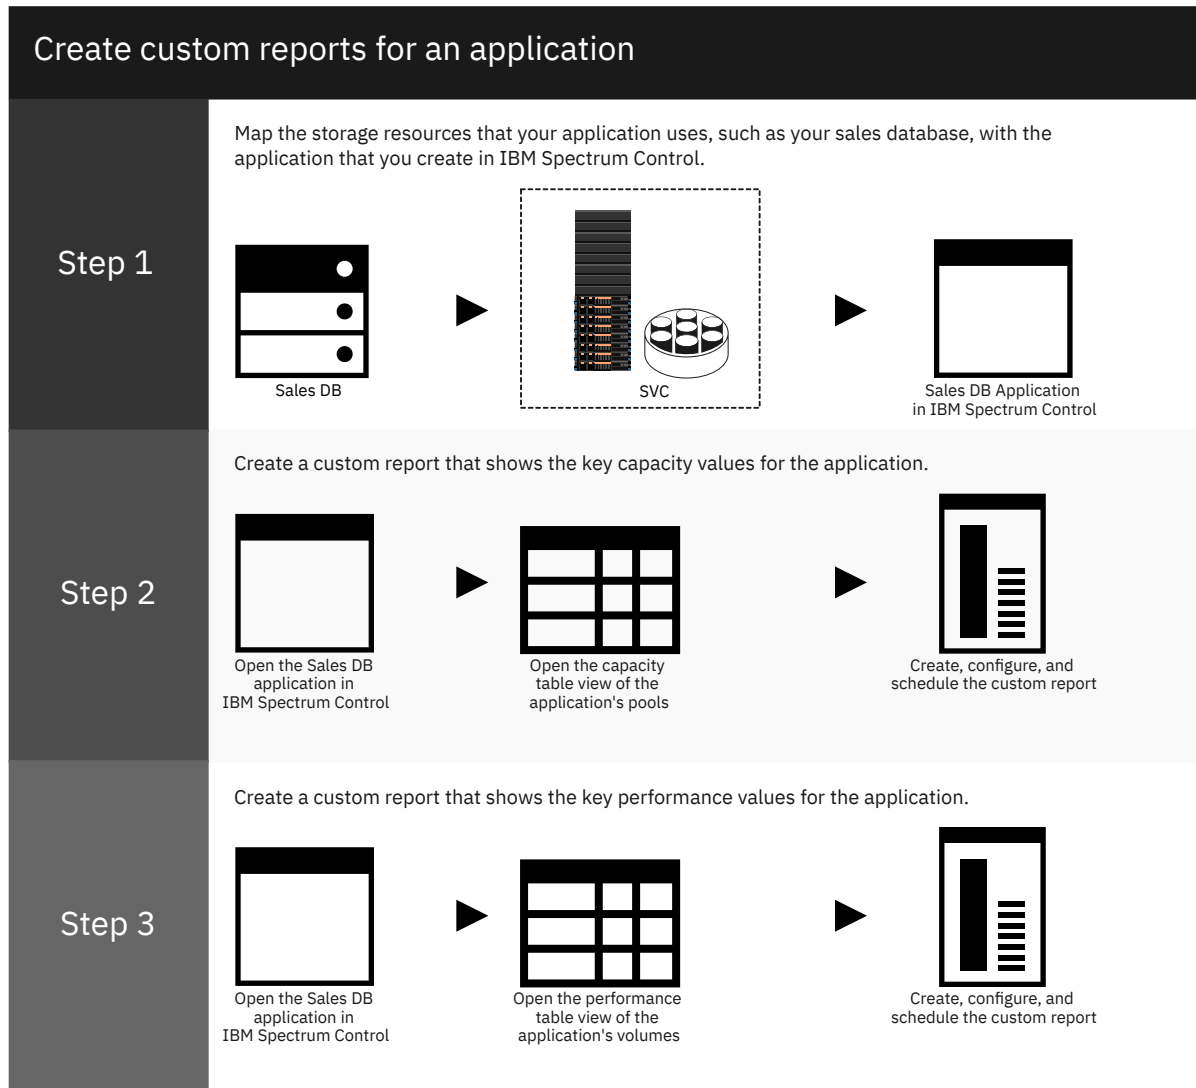
Related tasks

- [Configuring the email server](#)

Tutorials: Creating custom capacity and performance reports for applications

In these tutorials, you create an application so that you can create, schedule, and send reports by email about the storage resources that the application consumes.

About this task



You create the application IBM Spectrum® Control, and then create custom capacity and performance reports.

1. [Creating the application for the capacity report](#)
In this scenario, you create the application that you want to use to generate the custom report.
2. [Creating the capacity report for the application](#)
In this scenario, you create a scheduled report about the capacity and storage usage of an application's pools.
3. [Creating the performance report for the application](#)
In this scenario, you create a performance report about the volumes that are assigned to the application.

Creating the application for the capacity report

In this scenario, you create the application that you want to use to generate the custom report.

About this task

Add the storage resources that your application consumes to the application that you create in IBM Spectrum® Control.

Procedure

1. From the Groups menu, click Applications.
2. Click Create Application.
3. Complete these actions:
 - Enter the name.
 - Add a description.
 - Enter the type. For example, to distinguish this application from the other types of application that you create, enter DB or the type of DB, such as db2.
 - Enter the subtype. For example, you can enter preproduction, test, or, production to categorize applications by environment.
4. Click Assign storage resources to the application, and then click A selection of volumes that match a a pattern.
5. In this scenario, you enter an asterisk * as the name pattern for the volumes, then you click Belonging to, and add a name pattern that matches the names of the pools that contain the volumes.
Assigning resources: Enter name patterns instead of adding comma-separated lists of the resources that an application uses. When you add more storage resources that match the name pattern, they are dynamically added to the application.
6. Click Preview to check that you added the storage resources that the application consumes, and then click Save.

Next topic: [Creating the capacity report for the application](#)

Creating the capacity report for the application

In this scenario, you create a scheduled report about the capacity and storage usage of an application's pools.

About this task

You want to create and schedule a report about the capacity of the pools that are assigned to an application.

Procedure

1. From the Groups menu, click Applications.
2. In the Related Resources section, double-click the application, and then click Pools.
3. Click Create Report, enter the unique name of the report, and then click Next.
4. Select the capacity values that you want to include in the report.

In this scenario, the following actions are completed to configure the report output:

- a. In the General section, exclude all of the columns except for the Name and Tier columns.
- b. In the Storage section, add capacity values, and select Zero Capacity from the Other section.

What's zero capacity: The capacity information that is collected over 180 days is analyzed to determine, based on historical storage consumption, when the pools will run out of capacity. The pools that have already run out of capacity are marked as depleted. For the other pools, a date is provided so that you know when the pools are projected to run out of capacity. If sufficient information isn't collected to analyze the storage usage of the pool, `None` is shown as the value for zero capacity. If a capacity limit is set for the pool, the date shown in the Zero Capacity column is the date when the available capacity based on the capacity limit will be depleted. For example, if the capacity limit for a 100 GiB pool is 80%, it is the date when the available capacity of the pool is less than 20 GiB. `Depleted` is shown in the column when the capacity limit is reached.

- c. To reorder the report output, drag the column headings.

5. Add up to four filters.

Sample: Filter by Zero Capacity



Only include capacity information about the pools that are depleted or that could be analyzed to determine when the pools are projected to run out of capacity.

Tip: Don't forget to click Apply Filter when you add each filter, and check the report preview to ensure that the report's recipients will get the information that they need.

6. Choose one of the following scheduling options:

Option	Description
Email	Save and send the email now. You can also attach the report to the email as a CSV, PDF, or HTML file.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV, PDF, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

7. Click Save and Schedule.

Results

Depending on the scheduling option that you chose, the report is saved to your file system, sent, or will be sent at the interval that you specified. To preview or edit the report, click Reports>Reports, expand the Custom Report section, and select the report.

Previous topic: [Creating the application for the capacity report](#)

Next topic: [Creating the performance report for the application](#)

Creating the performance report for the application

In this scenario, you create a performance report about the volumes that are assigned to the application.

Procedure

1. Click Groups > Applications.
2. Double-click the application, and then click Volumes in the navigation pane.
3. Select the volumes and click View Performance.
Tip: Before you create the performance report, you can add more or other performance metrics to the table. The average, minimum, and maximum values of the performance metrics that are added to the table are calculated and shown in the report. paneClick Edit Table Metrics and select the additional metrics that you want in the report.
4. Click Create Report.
5. Enter the name of the report and click Next.
6. Select the metrics that you want to include in the report and click Next.
7. Add a filter and then click Apply Filter.
8. Choose one of the following scheduling options:

Option	Description
Email	Save and send the email now. You can also attach the report to the email as a CSV, PDF, or HTML file.
Save to File System	Schedule and save the report to the default reports folder or to the folder that you created for saving reports.
Email and Save to File System	Schedule the report, send it by email or attach the report to the email as a CSV, PDF, or HTML file, and save the report to the default reports folder or to the folder that you created for saving reports.

Results

The report is created. To edit the report, click Reports > Reports, expand the Custom Report section, and click the report.

Previous topic: [Creating the capacity report for the application](#)

Custom capacity and performance view reports

When you create custom capacity or performance view reports, you can specify a relative time range for the information that is shown about the resources in the reports.

To generate custom reports with relative time ranges, you must select one or more resources and click View Capacity, or View Performance. Alternatively, you can open the page for the resource and click the Capacity or Performance tab.

Remember: If you decide to specify your own start date, or end date, or both, you'll get the same information in the report each time it is generated. For example, you change the time range to 2 days, January 1 to January 2, and then generate a weekly report. Each time the report is generated, the information that was collected from January 1 to January 2 is used to generate the report.


Relative time ranges for custom capacity reports

When you create custom capacity reports on Capacity View pages, the default time range for the capacity information that is shown in the report comprises an aggregate of the capacity information that was collected over the previous month.


Choosing relative time ranges in custom capacity reports

Overview

When you create custom reports on the Capacity tab, you choose a relative time range. The time range determines how many daily collections of capacity information are analyzed when the report is generated.



Scenario



On November 18, you select the last week as the relative time range. You create the report and schedule the report to run on the fifth day of each month.

When the report is run on 5 December, the aggregate of the daily capacity information that was collected in the last week of November is shown in the report.

If you navigate to the capacity view for a storage resource, you can specify a relative time range for the capacity information in the report, such as last week, month, 6 months, or year. Depending on the time range that you specify, the aggregated values for the capacity information are shown in the report.


Relative time ranges for custom performance reports

When you create custom performance reports on Performance View pages, the default time range for the performance information that is shown in the report comprises an aggregate of the performance information that was collected over the last 12 hours.


Choosing relative time ranges in custom performance reports

Overview

When you create custom reports on the Performance tab, you choose a relative time range. The time range determines how many hourly or daily collections of performance information are analyzed when the report is generated.



Scenario



On November 18, you select the last week as the relative time range. You create the report and schedule the report to run on the fifth day of each month.

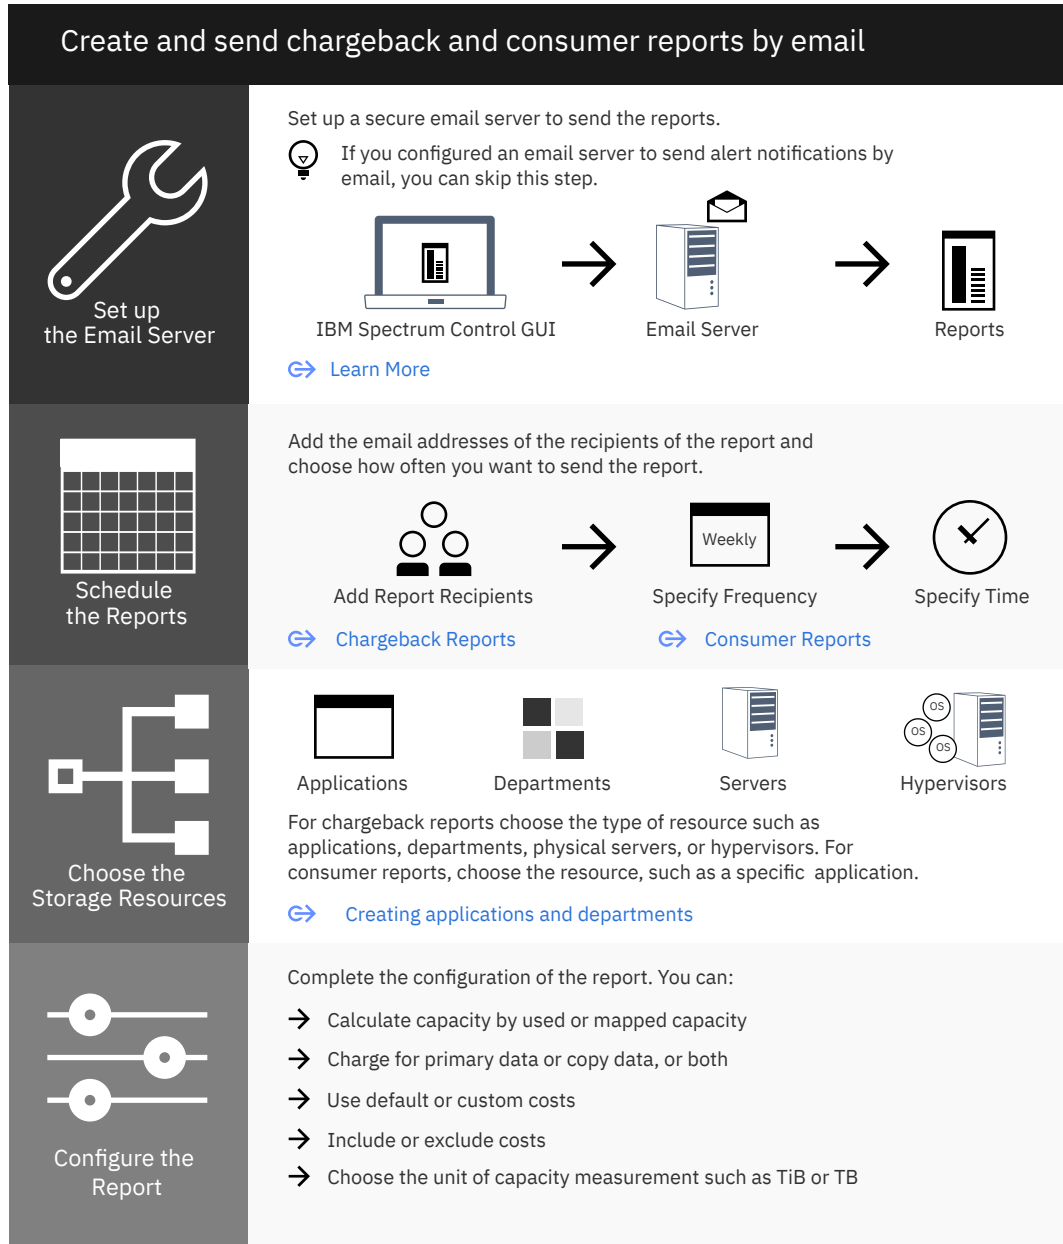
When the report is run on 5 December, the aggregate of the daily performance information that was collected in the last week of November is shown in the report.

If you navigate to the performance view for a storage resource, you can specify a relative time range for the capacity information in the report, such as the last hour, 6 hours, 12 hours, day, week, month, 6 months, or year. Depending on the time range that you specify, the aggregated values for the performance information are shown in the report.

Creating chargeback and consumer reports

To help plan capacity purchases and make your organization aware of the cost and the amount of the storage that is used by storage consumers, create chargeback and consumer reports.

Chargeback and storage consumer reports



In chargeback reports, you create reports that show the capacity and the cost of the block and file storage that is used by the types of storage consumers in your storage environment such as:

- Applications
- Departments
- Hypervisors
- Physical servers

In storage consumer reports, you create reports that show the capacity and cost of the block storage that is used by a single storage consumer, such as the capacity and the cost of the storage that is used by a single application.

To make the owners of the storage consumers aware of the capacity and cost of the storage they use, you schedule and send chargeback and storage consumer reports by email.

Show cost and capacity of tiered storage: If you want to show the block capacity of tiered storage and differentiate between the cost of storing data on tier 1 or on lower tiers of storage, you must assign your block storage pools to tiers. Click **Storage > Pools**. Right-click one or more pools, and then click **Tier** and the tier level from the menu. To view the capacity and charge for the storage that is used by applications and departments, you must create models of the applications and departments in your business organization and map the storage resources to the applications and map the applications that the departments use to the departments. If you organize an application or department in hierarchies, the chargeback and consumer reports show the capacity and cost of the storage resources that the application and its child applications use and the capacity and cost of the storage resources that the department and its child departments use.

Storage capacity

You can choose the unit of measurement that is used to show the amount of storage that is consumed by your storage resources. In consumer and chargeback reports, you can show the storage that is consumed, for example, in PB, PIB, TB, TiB, GB, or GiB.

Storage costs

The cost of the storage that is used by storage consumers is included in all reports even if the storage costs are set to zero. If you want to send reports that show the capacity of the storage, but not the cost of the storage that is used, you can exclude the storage costs from the report.

Before you decide how you want to set the costs for the block storage that is used by the storage consumers, you must choose whether you want to charge for the storage that is allocated to the storage consumer or the storage that is assigned to the storage consumer.

By default, block storage is calculated by the amount of storage that is allocated to the volumes that are being used for storage by the consumer. Whether you choose to calculate by allocated or assigned capacity, it doesn't affect the calculation of capacity for standard-provisioned volumes because standard-provisioned volumes are fully allocated when they are created. However, the calculation of the capacity of thin provisioned and compressed volumes is affected depending on whether you choose to calculate capacity by allocated or assigned capacity.

For example, an application has three volumes, two fully allocated volumes with 1 TiB each and one thin-provisioned volume with an assigned capacity of 3 TiB. The thin-provisioned volume is allocated capacity in increments of 1 TiB. If you generate a report that calculates capacity by allocated capacity, and the thin-provisioned volume is currently allocated 1 TiB, the total amount of capacity for the application is $1 \text{ TiB} + 1 \text{ TiB} + 1 \text{ TiB} = 3 \text{ TiB}$. However, if you choose to calculate by assigned space, the total capacity for the application is $1 \text{ TiB} + 1 \text{ TiB} + 3 \text{ TiB} = 5 \text{ TiB}$. The cost is also affected because the cost is the total capacity of the block storage, which is multiplied by the unit cost or the unit cost per tier.

Restriction: If you choose to calculate block storage by used capacity, the block storage for FlashSystem A9000 and FlashSystem A9000R storage systems is calculated based on the provisioned capacity percentage of the volumes. This restriction occurs because used capacity information for volumes is not available from FlashSystem A9000 and FlashSystem A9000R.

Default or custom storage costs

When you create or update a report, you can choose one of the following options:

Use default cost

Set or use default costs per unit and type of storage.

If you choose to use the default costs option, the changes that you make apply to all of the reports that use the default option to calculate the cost of storage. When you exclude costs from reports that use default costs, you exclude costs from all reports that use default costs to calculate the cost of storage.

Set custom cost

Use custom costs per unit and type of storage.

If you choose to use the custom costs option, the costs that you set are used to calculate the cost of storage only for that report. If you exclude costs from the custom report, the costs are only excluded from the custom report that you create or edit.

The reports that were created in IBM Spectrum® Control 5.2.11 use the unit costs that were specified when the reports were configured.

Tip: If you want to change the configuration of a report, such as how storage costs are calculated, click **Reports > Email Reports**, select the report, and then click **Edit**.

Deleted and unavailable reports

If you delete a storage consumer, such as an application, the entry for that storage consumer is removed from all chargeback reports for applications. If you created a consumer report for a storage consumer that was deleted, the report is automatically deleted.

If you organize your applications or departments in hierarchies, the chargeback and consumer reports show the total capacity and costs for the parent application or department and its child applications and departments. However, if you create a consumer report for a specific application or department and then designate the application or department as a child application of another application, the consumer report is deleted.

For example, you create a consumer report for an application called German Online Sales. You then reorganize your sales applications in a hierarchy and the application for German Online Sales becomes a child application of European Sales. To include the capacity and cost of the storage that is used for German Online sales, you must create a consumer report for the application for European Sales.

- **[Creating chargeback reports](#)**
Create chargeback reports that show how much capacity is used and the cost of the capacity that is used by applications, departments, hypervisors, and physical servers.
- **[Creating consumer reports](#)**
Create consumer reports that show how much block capacity is used and the cost of the block capacity that is used by an application, department, hypervisor, or physical server.
- **[Creating summary reports of the storage capacity](#)**
Configure reports that make managers aware of the capacity of the storage that is used by their applications, departments, hypervisors, or physical servers. The cost of the storage is not shown in the reports.

Creating chargeback reports

Create chargeback reports that show how much capacity is used and the cost of the capacity that is used by applications, departments, hypervisors, and physical servers.

About this task

Chargeback reports help managers to realize the cost of the storage that they use, plan storage usage more efficiently, and reduce storage costs. For example, a manager gets a chargeback report that shows that the bulk of the storage that is used by a non-critical application is on tier-1. To reduce storage costs and use storage more efficiently, the manager can request the storage administrator to use tier-2 or lower tiers of storage for the application's data.

Procedure

1. Click Reports, and then click Create Report.
2. In the Summary Reports pane, click Chargeback.
3. Type the unique name of the report, and enter the delivery and the scheduling details.
4. Choose the type of storage consumer from the list.
If you choose Hypervisor or Physical Server, the capacity and cost of the hypervisors and physical servers that belong to clusters are listed under the name of the cluster.
5. Choose how you want to calculate the capacity and cost of thin-provisioned storage.
6. Choose the unit of capacity measurement that you want to use in the report.
Tip: To make sure that the recipients of the report get the report that they want, check the report preview.
7. Choose whether you want to use default or customized costs for block and file storage.
To change the default costs for block and file storage, click the lock (🔒) icon.
8. By default, the storage costs are shown in the report. To exclude storage costs from the report, choose one of the following options:

Storage cost option	Steps
Use Default Cost	a. Click the lock (🔒) icon. b. In the Include Storage Cost in Report section, click No. Note: If you exclude costs from a report that uses default costs, the costs are excluded from all of the chargeback and consumer reports that use default costs.
Use Custom Cost	In the Include Storage Cost in Report section, click No.
9. Choose the types of storage that you want to include in the report.
10. Set the unit costs for the storage that is used.
To get a complete picture of the capacity and cost of the block storage, include the cost of the data that is copied.
11. Click Save.
If you set the frequency of the report to Now, you click Save and Send.

Results

The report is saved and, depending on the scheduling option, the report is sent to the recipients or scheduled to be sent later. If you want to change the configuration of the report, click Reports, select the report, and then click Edit.
Tip: You want to send a scheduled report now, but you don't want to change the original schedule? Click Reports, > Reports. Select the report, and click Actions, > Run Now. The report is sent without changing the original schedule.

What to do next

If you want to extract capacity information regularly about your storage resources, you can use the Representational State Transfer (REST) API to automatically export the data that you need to generate custom reports or share with external applications.

- [Chargeback reports](#)
Configure the chargeback reports to include block and file storage and to calculate the cost of the block and file storage that is used by the storage consumers.

Related information

- [REST API](#)

Chargeback reports

Configure the chargeback reports to include block and file storage and to calculate the cost of the block and file storage that is used by the storage consumers.

You configure, schedule, and send chargeback reports so that the owners and managers of storage resources:

- Know the cost of maintaining the block data and file data that they use
- Know the cost of maintaining block data on each tier of storage
- Know the cost of maintaining block copy data

Include total block and file capacity in the report

In the Configure Report section, you choose the types of storage that you want to show in the report. For example, if you click the File Storage check box, the names of the resources and the total file capacity that the resources use are shown in the report.

For block storage, you can show the total storage capacity and, if your storage is tiered, the total capacity by tier. You can also break down the total block capacity into total primary capacity and total copy capacity.

Tip: The recipient of the report, such as the application owner, sees that the capacity for copy data is high for tier-1 storage. To reduce costs and use the block storage more efficiently, the application owner can ask the storage administrator to place the volumes that are used for copy data on lower and less expensive storage tiers. The

storage administrator can also analyze the tiering of the pools to make sure that the application's performance is not affected when the copy data volumes are moved to lower storage tiers.

Include block and file costs in the report

In the Set Storage Cost section, you set the cost per TiB of storage that the resources use. The total cost of the storage that is used is shown in the Cost column in the report.

For block storage, you set the costs as follows:

- If your storage environment is tiered, you set the price per TiB for each tier of storage
- If your storage environment is not tiered, you set the price per TiB for all of your block storage
- If your storage environment contains tiered and non-tiered storage, you can set the price per TiB for the tiered storage, but you can't set a price for the non-tiered storage. Although you can't set the cost for the non-tiered storage, the capacity of the non-tiered storage that is used is calculated.

If you assign your block storage pools to tiers, you can differentiate between the cost of expensive tier-1 storage and the cost of the less expensive storage that is used on the lower tiers of storage.

For file storage, you set a price for all of the file storage that is used by your resources.

When you configure chargeback reports, you can send the report immediately or schedule the report to be sent every day, every week, or every month.

Tip: If your storage environment changes, you should check whether you need to update your scheduled reports. For example, if you add a storage tier, you can edit the report and set the price for the new tier.

Creating consumer reports

Create consumer reports that show how much block capacity is used and the cost of the block capacity that is used by an application, department, hypervisor, or physical server.

About this task

Storage consumer reports help the managers and owners of applications, departments, hypervisors, and physical servers realize the cost of the block storage that they use, plan block storage usage more efficiently, and reduce block storage costs.

Procedure

1. Click Reports, and then click Create Report.
2. In the Summary Reports pane, click Storage Consumer.
3. Type the unique name of the report, and enter the delivery and the scheduling details.
4. From the list, click the type of storage consumer and type or click the name of the storage consumer.
To generate a report about the hypervisors or physical servers that belong to a cluster, you choose Hypervisor or Physical Server, and then choose or type the name of the cluster.
5. Choose whether you want to calculate block capacity by allocated or assigned space.
6. Choose the unit of capacity measurement that you want to use in the report.
Tip: To make sure that the recipients of the report get the report that they want, check the report preview.
7. Choose whether you want to use default or customized costs for block storage.
8. Choose the types of storage that you want to include in the report.
To change the default costs for block storage, click the lock (🔒) icon.
9. By default, the storage costs are shown in the report. To exclude storage costs from the report, choose one of the following options:

Storage cost option	Steps
Use Default Cost	a. Click the lock (🔒) icon. b. In the Include Storage Cost in Report section, click No. Note: If you exclude costs from a report that uses default costs, the costs are excluded from all of the chargeback and consumer reports that use default costs.
Use Custom Cost	In the Include Storage Cost in Report section, click No.

10. Set the unit costs for the block storage that is used.
To get a complete picture of the capacity and cost of the block storage, include the cost of the data that is copied.
11. Click Save.
If you set the frequency of the report to Now, you click Save and Send.

Results

The report is saved and, depending on the scheduling option, the report is sent to the recipients or scheduled to be sent later. If you want to change the configuration of the report, click Reports, select the report, and then click Edit.

Tip: You want to send a scheduled report now, but you don't want to change the original schedule? Click Reports > Reports. Select the report, and click Actions > Run Now. The report is sent without changing the original schedule.

What to do next

If you want to extract capacity information regularly about your storage resources, you can use the Representational State Transfer (REST) API to automatically export the data that you need to generate custom reports or share with external applications.

- [Consumer reports](#)
Configure consumer reports to calculate the cost and capacity of the block storage that is used by a storage consumer such as an application, a department, a

hypervisor, or a physical server.

Related information

- [REST API](#)

Consumer reports

Configure consumer reports to calculate the cost and capacity of the block storage that is used by a storage consumer such as an application, a department, a hypervisor, or a physical server.

You configure, schedule, and send consumer reports so that the owner or manager of the storage consumer knows the cost of maintaining the block storage that the storage consumer uses.

In the Configure Report section, you include the total block storage capacity and, if your storage is tiered, the total capacity by tier in your report. You can also include the total primary capacity and the total copy capacity of the storage that is used, and you can include the total capacity:

- For VDisk mirrors
- For FlashCopy® and Safeguarded Copy volumes
- For remote mirrors such as Global Mirror, Global Copy, and Metro Mirror

In the Set Storage Cost section, you set the cost per TiB of the block storage that the storage consumer uses.

You set the costs as follows:

- If your storage environment is tiered, you set the price per TiB for each tier of storage
- If your storage environment is not tiered, you set the price per TiB for all of your block storage
- If your storage environment contains tiered and non-tiered storage, you can set the price per TiB for the tiered storage, but you can't set a price for the non-tiered storage. Although you can't set the cost for the non-tiered storage, the capacity of the non-tiered storage that is used is calculated.

Creating summary reports of the storage capacity

Configure reports that make managers aware of the capacity of the storage that is used by their applications, departments, hypervisors, or physical servers. The cost of the storage is not shown in the reports.

About this task

You can create reports that show the total capacity of the storage that is used by storage consumers and exclude the storage costs from the reports.

Procedure

1. Click Reports > Reports, and then click Create Report.

2. Choose one of the following options:

Option	Steps
Create chargeback reports	In the Summary Reports pane, click Chargeback.
Create consumer reports	In the Summary Reports pane, click Storage Consumer.

3. Enter the report details and schedule the report.

4. Choose one of the following options:

Option	Steps
Chargeback reports	Choose the type of storage consumer.
Consumer reports	Choose the type of storage consumer and the storage consumer.

5. To calculate the capacity that is assigned to the thin-provisioned volumes, click Assigned space.

By default, the amount of space that is allocated to the thin-provisioned volumes is used to calculate the cost.

6. Choose the unit of capacity measurement that you want to use in the report.

Tip: To make sure that the recipients of the report get the report that they want, check the report preview.

7. In the Configure Report section, choose the capacity values that you want to include in the report.

8. By default, the storage costs are shown in the report. To exclude storage costs from the report, choose one of the following options.

If you exclude costs from a report that uses default costs, the costs are excluded from all of the chargeback and consumer reports that use default costs.

Storage cost option	Steps
Use Default Cost	a. Click the lock (🔒) icon. b. In the Include Storage Cost in Report section, click No.
Use Custom Cost	In the Include Storage Cost in Report section, click No.

9. Click Save.

If you set the frequency of the report to Now, you click Save and Send.

Results

The reports include the capacity values and exclude the costs of the capacity that is used.

Running reports

You can run reports that you created without defining a schedule or changing the schedule that was created for the report.

About this task

If you created a report, but didn't schedule it, you can run the report again. You can also run reports that you scheduled without changing the schedule for the report.

Procedure

1. Click Reports.
2. Click the report.
To run two or more reports, press Ctrl or press Shift and click the reports.
3. Click Actions > Run Now.

Results

The reports are generated and sent to the reports' recipients.

Editing reports

Change the configuration and scheduling for reports.

About this task

You can edit and reschedule custom, predefined capacity and inventory, and consumer and chargeback reports.

Procedure

1. From the Reports menu, click Reports.
2. Expand the section that contains the report.
3. Click the report, and then click Edit.
Tip: For custom and predefined capacity and inventory reports, click one of the links in the task pane. By default, the first page that is shown is the Schedule Delivery page. For example, to add, change, or remove filters, click Edit Filter.
4. Make your changes.

Results

If you schedule the report, the changes are saved and included in the next scheduled report. If you chose send now when you scheduled the report, the changes are saved and the report is sent to the recipients.

What to do next

Don't forget when you send the report that you will need to reconfigure the schedule if you want to send or save the report at regular intervals.

Deleting reports

Remove the reports that you don't need.

Procedure

1. Click Reports.
2. Click the report.
To select two or more reports, press Ctrl or press Shift and click the reports.
3. Click Actions > Delete Report.

Configuring the email server

To send reports, you must configure the email server.

About this task

The email server that is used to send alert notifications is also used to send reports. If you didn't configure the email server for alert notifications, you must configure the email server before you create reports.

IBM Spectrum® Control ensures, when it sends alert notifications and reports by email to your email server, that it complies with the security standards that you configure on your email server. For example, if your email server requires authentication and TLS to send emails, then IBM Spectrum Control uses the authentication credentials that you provide when you set up the email server, and uses TLS to establish the connection with your email server. Because the email server that you set up is automatically trusted, you don't have any additional management tasks such as importing the server certificate into the keystores that are used by IBM Spectrum Control.

Tip: To change the configuration of the mail server, you click [Settings](#) > [Notification Settings](#).

Procedure

1. Type the host name, or IPv4, or IPv6 address of the email server.
2. Type the port number that is used by the email server to send alert notifications and reports.
3. Provide the user's credentials for authenticating with your email server.
If your email server requires authentication to establish the connection, you must enter the user name and password.
4. Optional: Type the email address, such as the email address of the administrator, that you want to use for receiving replies to alert notifications and reports.
If you don't enter a value, and you entered an email address in the User Name field, then replies are sent to that address. Otherwise, replies are sent to **no-reply@<Spectrum_Control_server_host_name>**.
5. Test the connection and then click Save.

Results

You can create reports and send the reports by email.

Investigating issues with reports

You can access log and trace files to investigate the issues that you might have when reports are generated and sent by email or saved to the file system.

In IBM Spectrum® Control, you can create, schedule, and send chargeback and storage consumption reports by email, and you can create, schedule, and save custom reports to your file system or send them by email.

To create reports, the following IBM Spectrum Control components and services are used. (The Export server service is used only for custom reports.)

Table 1. Components used in reporting

Actions	Web server component	Data server component	Device server component	Export server service	SMTP transmission service
Create, update, and delete reports	Yes	Yes	Yes	No	No
Generate and deliver reports	Yes	Yes	No	Yes	Yes

Export server is a new service that converts the data for the report that it gets from the Web server into HTML.

The SMTP transmission service handles all communications, such as the delivery of alert notifications by email, and the sending of chargeback, storage consumption, and custom reports from IBM Spectrum Control to the customers' external SMTP servers.

Setting the trace level for investigating reporting issues

Go to the folder where you installed IBM Spectrum Control and complete these steps:

Tip: For Windows operating systems the default installation directory is C:\Program Files\IBM\TPC. For AIX® and Linux® operating systems the default installation directory is opt/IBM/TPC.

1. On Windows OS, go to installation_dir\IBM\TPC\web\conf. On AIX or Linux OS, go to installation_dir/IBM/TPC/web/conf.
2. In a text editor, open the logging.properties file.
3. Set the value of the **trace.reportdata.level** property to **trace.reportdata.level=debug**.
4. Save the logging.properties file.

You have set the trace level for the traceReportData_<number>.log to debug.

Access log, trace, and other information about reports

Table 2. Location of log, trace, and other files

File name	Description and location
ReportData.log	High-level information about the creation and distribution of reports. Windows: installation_dir\web\log AIX and Linux: installation_dir/web/log
traceReportData_<number>.log	Detailed information about the creation and distribution of custom reports, and SMTP transmission of custom reports. For example, you can check why a custom report isn't saved to the file system, or sent, or received by email. Windows: installation_dir\web\log AIX and Linux: installation_dir/web/log
traceWebServer_<number>.log	Information relating to creating, editing, and scheduling custom reports. Windows: installation_dir\web\log AIX and Linux: installation_dir/web/log
Console.log	Information about whether the report data service, which runs on the Web server, was started or stopped, errors generated when creating and retrieving custom reports, handshake errors that were generated when custom reports are sent by email. Windows: installation_dir\wlp\usr\servers\webServer\logs AIX and Linux: installation_dir/wlp/usr/servers/webServer/logs

File name	Description and location
Messages.log	See the description for the Console.log. Both log files report the same types of errors. Windows: installation_dir\wlp\usr\servers\webServer\logs AIX and Linux: installation_dir/wlp/usr/servers/webServer/logs
traceDeviceServer_<number>.log	Detailed information about the SMTP transmission of the chargeback and storage consumption reports sent by email. For example, you can check why a chargeback or storage consumption report wasn't sent. Windows: installation_dir\device\log AIX and Linux: installation_dir/device/log
Messages.log	Information about whether the Device server was started or stopped, handshake errors that are generated when alert notifications, chargeback, or storage consumption reports are sent. Windows: installation_dir\device\log AIX and Linux: installation_dir/device/log
traceScheduler_<number>.log	Information about when the report creation process began and completed for each report. Windows: installation_dir\data\log AIX and Linux: installation_dir/data/log

Stopping and starting the Export server service

On Windows, the Export server runs as a service called IBM Spectrum Control - Export Server, which can be stopped and restarted.

1. Click the Start menu, type `services.msc`, and then press Enter.
2. Click IBM Spectrum Control - Export Server, and stop or start it.

Alternatively, you can run the following batch files or scripts to stop or start the Export server service on Windows or AIX and Linux operating systems:

Operating system	Location	File
Windows	By default, it's C:\Program Files\IBM\TPC\scripts	Start: startTPCExport.bat Stop: stopTPCExport.bat
AIX and Linux	By default it's /Opt/IBM/TPC/scripts	Start: startTPCExport.sh Stop: stopTPCExport.sh

Types of predefined capacity and inventory reports

The predefined capacity and inventory reports that you can create for your storage resources are listed.

Check the table to see the predefined reports that you can create for the storage resources in your environment.

Table 1. Predefined capacity and inventory reports

Resource	Capacity	Inventory
Block Storage Systems	Yes	Yes
File Storage Systems	Yes	Yes
Object Storage Systems	Yes	Yes
Switches	No	Yes
Chassis	No	Yes
Servers	Yes	Yes
Hypervisors	Yes	Yes
Applications	Yes	No
Departments	Yes	No
Managed Disks	Yes	Yes
Controllers	No	Yes
Enclosures	No	Yes
Filesets	Yes	No
NAS File Systems	Yes	No
Modules	No	Yes
Network Shared Disks	Yes	No
Pools	Yes	No
Spectrum Scale Nodes	No	Yes
Spectrum Virtualize Nodes	No	Yes
Storage System FC Ports	No	Yes
Storage System IP Ports	No	Yes
Server Ports	No	Yes
Switch Ports	No	Yes
Volumes	Yes	No

Reports FAQ

Troubleshooting reports

Why can't I send the report as an attachment?

You're notified that a report can't be sent by email because the report's attachments exceed the maximum size limit. The maximum size limit for attachments is set on your email server.

If you can't change the maximum size limit for attachments, you can edit the report and choose Save to File System as the delivery method on the Schedule Delivery page.

Alternatively, you can reduce the scope of the report to create smaller and more focused reports. For example, when you create reports about volumes, you can select volumes by the storage systems, pools, or applications that they belong to, or the servers that they are assigned to. You can also refine the report by removing columns and adding filters. And, to reduce the size of attachments, choose CSV as the file type.

[Learn more](#)

Why can't I save the report to a custom folder on my file system?

You must have read and write access to the folder and you must create the custom folder before you create the report. Don't forget to enter the full path to the folder, such as C:\Program files\IBM\TPC\myreports or /opt/IBM/TPC/myreports on the Schedule Delivery page.

Why can't I see all of the rows in the report in my email?

If the report is very large, the report's recipients might not see all of the rows in the report in the email.

Click Reports. Select the report and click Edit. On the Schedule Delivery page, add the report as a CSV or HTML attachment. To see all of the rows in the report, the reports' recipients can open the attachment.

Creating and sending

Why can't I create a new custom report on the Reports page?

Unlike chargeback, consumer, and predefined reports, you can't create new custom reports on the Reports page. To create new custom reports, open any page that shows asset, capacity, configuration, or performance information, such as the Block Storage Systems page, and click Create Report. On the Reports page, you can edit or delete the custom reports that you created.

[Learn more](#)

How do I create a custom report about performance?

Go to a page that shows information about block storage resources. Select one or more of the resources, click View Performance, and then click Create Report.

[Learn more](#)

What types of reports can I create?

IBM® Storage Insights Pro users can create chargeback, consumer, predefined capacity, predefined inventory, and custom reports. IBM Storage Insights users can create predefined capacity reports for storage systems and pools and predefined inventory reports for storage systems. To create new predefined reports, click Reports, and then click Create Report.

[Learn more](#)

How can I send the reports that I created now?

You want to send a report now although the report, for example, is scheduled to run next week. Or, you created a report that you didn't schedule and want to send again.

Click Reports. To select the reports, press Ctrl or press Ctrl + Shift and click the reports. From the Actions menu, click Run Now. The report is generated without changing the schedule that you created for the report.

Capacity metrics for reports

To gain insights into storage usage, review the capacity metrics that are collected, analyzed, and shown in reports.

- [Capacity metrics for chargeback and consumer reports](#)
Learn more about how the capacity and costs for block and file storage is calculated in chargeback and consumer reports.

Capacity metrics for chargeback and consumer reports

Learn more about how the capacity and costs for block and file storage is calculated in chargeback and consumer reports.

Chargeback and storage consumer reports

The same block capacity metrics are shown in chargeback and consumer reports. In chargeback reports, the block capacity and costs are shown for all applications, departments, hypervisors, and physical servers. In consumer reports, the block capacity and costs are shown for a specific application, department, hypervisor, or physical server. The capacity and cost of file storage is not included in consumer reports.

If you create reports about physical servers, servers that are virtual machines are excluded from the report.

Tip: To find out which servers are virtual machines, click Servers > Servers, right-click any column heading and click Virtual Machine.

Virtual machine capacity

If you assign the storage resources on a virtual machine to an application or department, the capacity of those storage resources is included in the calculation of capacity and cost for the application or department.

For example, three virtual machines are allocated capacity from a data store, which has a total capacity of 3 TiB. The capacity for the data store comprises:

- One volume on tier 1 with a capacity of 2 TiB
- One volume on tier 2 with a capacity of 1 TiB

The ratio of tier-1 storage to tier-2 storage is 2:1 or 66.66% to 33.33%.

The total capacity of the data store is split across 3 by 1,024 GiB disks and one of the disks is allocated to `Virtual_Machine_1`. The used capacity on `Virtual_Machine_1` is allocated to a database application. So, the total amount of capacity that is allocated to the database application is 1,024 GiB, which comprises 66.66% of tier-1 storage and 33.33% of tier-2 storage. When the report is generated for the database application, the block capacity for tier-1 storage is **1024 GiB*66.66%**, which is 682.67 GiB or 0.67 TiB. The block capacity for the tier-2 storage is **1024 GiB*33.33%**, which is 341.33 GiB or 0.33 TiB.

If you calculate the capacity and cost of storage by used capacity, the capacity and cost of the storage that is used by the application or department is shown in the chargeback and consumer reports.

Restriction: You might get unexpected results in chargeback and consumer reports if the volumes that are used to provide the capacity on the disks for the virtual machine are thin-provisioned and shared by two or more applications.

Hypervisor and physical server clusters

In chargeback and storage consumer reports, the capacity and cost of storage for hypervisors and physical servers that are members of clusters are included in the capacity and cost of the cluster that they belong to. For example, `hypervisor_1` and `hypervisor_2` belong to `cluster_1`. To create the chargeback report about the hypervisors cluster, you choose Hypervisors as the consumer, and the capacity and the cost of the two hypervisors are included in the report entry for `cluster_1`. To create a report for `cluster_1` in a storage consumer report, you choose Hypervisor as the type of consumer, and `cluster_1` as the name of the consumer.

Block storage

The following information is provided about how block capacity is calculated:

Block or Used Capacity

In chargeback reports, the total block capacity for all of the storage resources is called `Block capacity`, whereas in consumer reports the total block capacity that the storage consumer uses is called `Used Capacity`.

The total capacity of these volumes is included in the calculation of block capacity:

- The provisioned capacity of the volumes without copies that are allocated to the consumer
- The provisioned capacity of the volumes with copies that are allocated to the consumer such as:
 - The total capacity of VDisk mirror volumes and their copies
 - The total capacity of FlashCopy® volumes and their target volumes
 - The total capacity of Safeguarded Copy source volumes and their volume backups
 - The total capacity of Global Mirror, Global Copy, and Metro Mirror volumes and their target volumes

For example, an application is assigned all of the block storage volumes on a physical server. In this scenario, the block capacity value includes the capacity of all of the block volumes that are assigned to the physical server. If the assigned volume on the physical server has a VDisk mirror copy or the volume is a FlashCopy source volume, the capacity of the VDisk mirror copy or the capacity of the FlashCopy target volume is also included in the calculation.

Primary capacity

The total capacity of the following types of volumes is included in the calculation of primary capacity:

- The provisioned capacity of the volumes without copies that are allocated to the consumer
- The provisioned capacity of primary VDisk mirror copies that are allocated to the consumer
- The provisioned capacity of the source volumes in FlashCopy relationships that are allocated to the consumer
- The provisioned capacity of the Safeguarded Copy source volumes that are allocated to the consumer
- The provisioned capacity of the Global Mirror, Global Copy, and Metro Mirror source volumes that are allocated to the consumer

VDisk Mirrors

The total capacity of the VDisk mirror copies for the volumes that are allocated to the consumer.

FlashCopy

The total capacity of the following types of volumes is included in the calculation of FlashCopy:

- The capacity of the FlashCopy target volumes for the source volumes that are allocated to the consumer
- The capacity of the volume backups for the Safeguarded Copy source volumes that are allocated to the consumer

Remote Mirrors

The total capacity of the Global Mirror, Global Copy, and Metro Mirror target volumes for the source volumes that are allocated to the consumer.

Primary capacity by tier

The primary capacity that is allocated to the consumer for the specified tier.


Copy capacity by tier

The copy capacity that is allocated to the consumer for the specified tier, which is calculated as follows:

- The total capacity of the non-primary VDisk mirror copies
- The total capacity of the target volumes for volumes in FlashCopy relationships
- The total capacity of target volumes for Global Mirror, Global Copy, and Metro Mirror

In tiered environments, the source volume might be on tier 1 and the target volume on tier 3. In this case, the capacity of the source volume is included in the calculation for the total primary capacity for Tier 1 and the capacity of the target volume is included in the calculation for the total copy capacity for tier 3.

Figure 1. Calculating copy capacity



Restriction: For volumes that have multiple volume copies, only the capacity of the first copy and the capacity of the second copy of the volume are included in the calculation of copy capacity.

File storage

File capacity is the total amount of file storage space that is allocated to the consumer.

For example, if a server mounts a fileset export from a file storage system and a quota is defined for the fileset, the file capacity is the hard limit capacity that is defined for the quota.

If a quota is not defined for the fileset, the file capacity is the capacity of the file system that contains the fileset. If multiple fileset exports from the same file system are mounted and the hard limits that are defined for the quota exceed the total file system capacity, then the file capacity is the total capacity of the file system.

If a hypervisor data store is a mounted NFS export, file capacity is the capacity of the file system on which the export is based.

Using the REST API to generate reports

Use the REST command line utility or a web browser to generate capacity, configuration, and performance reports.

- [Retrieving data about resources by using a REST API command line utility](#)
You can connect to the REST API for IBM Spectrum® Control and retrieve data by using a REST command line utility.
- [Retrieving data by using REST API with a web browser](#)
A convenient way to use the IBM Spectrum Control REST API is within the interface itself. You can quickly access resource information for your reporting capabilities by using a web browser.

Retrieving data about resources by using a REST API command line utility

You can connect to the REST API for IBM Spectrum® Control and retrieve data by using a REST command line utility.

You can use any REST command line utility, for example, GNU Wget. To access information about the GNU Wget utility, go to <https://www.gnu.org/software/wget/>. The IBM Spectrum Control REST API is hosted here: `https://<hostname>:9569/srm/REST/api/v1/`.

You can enter a command to authenticate with the REST API and store a security token in a file to use in your later queries, for example, `cookies.txt`.

You can use the utility, by entering the following command:

```
wget --post-data "j_username=<user name>&j_password=<password>"  
--no-check-certificate --keep-session-cookies --save-cookies cookies.txt  
https://<hostname>:9569/srm/j_security_check
```

You can use the security token in all subsequent commands that are issued against the REST API, by entering the following command:

```
wget --no-check-certificate --load-cookies cookies.txt  
https://<hostname>:9569/srm/REST/api/v1/
```

Example (partial listing)

```
},  
{  
  "Description": "Provides a list of ports belonging to a parent storage system.  
Format: \StorageSystems\<id>\Ports For a list of ports associated with a  
port owner, specify the owner type and id in the URL.
```

To see a list of all storage systems, you can append the URL:

```
https://<hostname>:9569/srm/REST/api/v1/StorageSystems
```

To see more information about a resource type, you can further append the URL (where `<Name>` is the resource type, such as `StorageSystems`, `Switches`, `Servers`, etc.)

```
https://<hostname>:9569/srm/REST/api/v1/<Name>
```

To see more information about a specific resource, add the ID to the URL (where `<id>` is available from `https://<hostname>:9569/srm/REST/api/v1/<Name>`):

```
https://<hostname>:9569/srm/REST/api/v1/<Name>/<id>
```

To query for volumes that belong to a specific storage system, enter:

```
https://<hostname>:9569/srm/REST/api/v1/StorageSystems/57909/Volumes
```

Example (partial listing)

```
[  
  {  
    "Acknowledged": "No",  
    "Allocated Space": "1.00",  
    "Capacity": "1.00",  
    "Controller": "Node 0",  
    "Copy Relationship": "",  
    "Easy Tier": "Tiered Pools\No",  
    "Encryption": "",  
    "Enterprise HDD Capacity": "1.00",  
    "Format": "FB",  
    "Hosts": "18",  
    "LSS or LCU": "00",
```

```

    "Last Data Collection": "Apr 10, 2016, 15:03:57",
    "Name": "CET_RA_vol",
    "Nearline HDD Capacity": "",
    "Physical Allocation": "Fully Allocated",
    "Pool": "General Use 1",
    "RAID Level": "RAID 5",
    "SSD Capacity": "",
    "Service Class": "",
    "Shortfall": "",
    "Status": "Normal",
    "Storage System": "DS8000-2107-75BLG91-IBM",
    "Storage Virtualizer": "None",
    "Thin Provisioned": "No",
    "Ticket": "",
    "Tier Distribution": "0",
    "Unallocated Space": "0.00",
    "Unused Space": "",
    "Used Allocated Space": "",
    "Used Space": "1.00",
    "Virtualizer Disk": "None",
    "Volume ID": "0000",
    "Volume Number": "0",
    "Volume Unique ID": "75blg91\0000",
    "id": "79852"
  },

```

The information that is provided in the **Format** field in the example shows what combinations are possible. For example,

https://<hostname>:9569/srm/REST/api/v1/StorageSystems/<id>/Ports

Tips:

- If the **Format** field is omitted, you can add the ID of one resource to the URL.
- The name elements in the URL are case-sensitive.
- The names of resource types are plural.
- If the URL is not used correctly, you might receive the following error message:

```

{ "result": { "type": "E", "msgId": "BPCUI0099E", "time": "Apr 4, 2016 16:25:07", "text":
"The storage resource is not available." } }

```

- If the URL returns a blank page, it means that no data is available for the resource. For example, if you query an application for file shares and there are no file shares in the application, the resulting page is blank.

Table 1. IBM Spectrum Control REST API URL examples

IBM Spectrum Control REST API	Description
https://<hostname>:9569/srm/REST/api/v1/StorageSystems	A list of all monitored storage systems.
https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002	Specific information for one storage system (ID is available from this URL: https://<hostname>:9569/srm/REST/api/v1/StorageSystems/<id>).
https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes	A list of all volumes for one storage system.
https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance	The available performance metrics for a volume.
https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/RemoteReplication	The remote mirror replication relationships for storage systems.
https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance/824?granularity=sample&startTime=1455818400000&endTime=1455904800000	Sample performance granularity information and a start and end time for the performance data.

Note: IBM® Storage is changing the capacity terminology that is used in IBM Storage products to make it more consistent. The new capacity terminology is used in IBM Spectrum Control 5.3.6 or later. The new terminology is used in the GUI, but the names used in the REST API have not changed. The output from the REST API continues to use Allocated Space instead of Used Capacity, Physical Allocation instead of Used Capacity %, and so on.

Retrieving data by using REST API with a web browser

A convenient way to use the IBM Spectrum® Control REST API is within the interface itself. You can quickly access resource information for your reporting capabilities by using a web browser.

To use the REST API with a web browser, use these steps:

1. Log on to the IBM Spectrum Control GUI. To start IBM Spectrum Control, see [Starting IBM Spectrum Control](#).
2. Open a new tab in your web browser.
3. Enter **https://<hostname>:9569/srm/REST/api/v1/**.

For example, to see a list of fabrics, enter **https://<hostname>:9569/srm/REST/api/v1/Fabrics** in a new tab of your web browser (partial listing).

```

{
  {
    "Acknowledged": "No",
    "Active Zone Set": "",
    "Connected Switch Ports": "0",
    "Custom Tag 1": "",
    "Custom Tag 2": "",
    "Custom Tag 3": "",
    "Data Source Count": "1",
    "Fabric Type": "Cisco",

```

```

    "Last Successful Probe": "N/A",
    "Links": "0",
    "Location": "",
    "NPV Switches": "0",
    "Name": "VSAN0002",
    "Parent Fabric": "san_director",
    "Principal Switch of Fabric": "tpc-70swt-csc",
    "Probe Schedule": "N/A",
    "Probe Status": "N/A",
    "Status": "Error",
    "Switch Ports": "0",
    "Switches": "1",
    "Virtual": "Yes",
    "WWN": "2002000DECAC5081",
    "id": "2663"
}

```

To see the information for a specific fabric in your environment, append the URL with the ID:

<https://<hostname>:9569/srm/REST/api/v1/Fabrics/2017>

Example (partial listing)

```

{
  "Acknowledged": "No",
  "Active Zone Set": "",
  "Connected Switch Ports": "18",
  "Custom Tag 1": "",
  "Custom Tag 2": "",
  "Custom Tag 3": "",
  "Data Source Count": "2",
  "Fabric Type": "Cisco",
  "Last Successful Probe": "N/A",
  "Links": "4",
  "Location": "",
  "NPV Switches": "0",
  "Name": "san_director",
  "Parent Fabric": "",
  "Principal Switch of Fabric": "",
  "Probe Schedule": "N/A",
  "Probe Status": "N/A",
  "Status": "Error",
  "Switch Ports": "104",
  "Switches": "2",
  "Virtual": "No",
  "WWN": "",
  "id": "2017"
}

```

To see the switches that belong to the fabric `"id": "2017"`, enter:

<https://<hostname>:9569/srm/REST/api/v1/Fabrics/2017/Switches>

Example (partial listing)

```

[
  {
    "Acknowledged": "No",
    "Connected Fabrics": "",
    "Custom Tag 1": "",
    "Custom Tag 2": "",
    "Custom Tag 3": "",
    "Data Source Count": "1",
    "Domain ID": "10",
    "Fabric": "unstable_BRCD_1-11",
    "Firmware": "v6.4.3g",
    "IP Address": "9.11.91.241",
    "Last Successful Monitor": "None",
    "Last Successful Probe": "None",
    "Links": "",
    "Location": "",
    "Mode": "Native",
    "Model": "Brocade 4100",
    "Name": "mdm-y76-swt",
    "Parent Switch": "",
    "Performance Monitor Interval (min)": "",
    "Performance Monitor Status": "Disabled",
    "Ports": "",
    "Principal Switch of Fabric": "cjswitch4",
    "Probe Status": "Never Probed",
    "Serial Number": "1070029",
    "Status": "Normal",
    "Vendor": "IBM",
    "Virtual": "No",
    "WWN": "100000051E347790",
    "id": "60989"
  }
]

```

Note: IBM Spectrum Control can connect to Brocade switches and fabrics either directly or by using Brocade Network Advisor. For Brocade switches and fabrics that IBM Spectrum Control connects to directly, the Data Source Count value is 0.

Table 1. IBM Spectrum Control REST API URL examples

IBM Spectrum Control REST API	Description
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems">https://<hostname>:9569/srm/REST/api/v1/StorageSystems	A list of all monitored storage systems.

IBM Spectrum Control REST API	Description
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002	Specific information for one storage system (ID is available from this URL: <a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems">https://<hostname>:9569/srm/REST/api/v1/StorageSystems).
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes	A list of all volumes for one storage system.
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance	The available performance metrics for a volume.
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/RemoteReplication">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/RemoteReplication	The remote mirror replication relationships for storage systems.
<a href="https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance/824?granularity=sample&startTime=1455818400000&endTime=1455904800000">https://<hostname>:9569/srm/REST/api/v1/StorageSystems/2002/Volumes/2465/Performance/824?granularity=sample&startTime=1455818400000&endTime=1455904800000	Sample performance granularity information and a start and end time for the performance data.

Note: IBM® Storage is changing the capacity terminology that is used in IBM Storage products to make it more consistent. The new capacity terminology is used in IBM Spectrum Control 5.3.6 or later. The new terminology is used in the GUI, but the names used in the REST API have not changed. The output from the REST API continues to use Allocated Space instead of Used Capacity, Physical Allocation instead of Used Capacity %, and so on.

Exporting information about resources

Notify your colleagues about the current state and potential issues with your storage environment by exporting the capacity and performance information that is shown in the GUI.

You can use the information that is shown about your storage resources and their related resources to create reports in HTML or PDF file formats that you can send to your colleagues. Alternatively, you can export the information about your resources to a CSV file, which you can include in a report that you want to create about your storage resources.

You can go to the resource or internal resource page and create reports that contain the following types of information:

Table 1. Types of information in resource reports

Location	Types of information
Resource or internal resource page	Status, configuration, and capacity data
Performance tab for resource or internal resource	Key performance metrics
Capacity tab for resource or internal resource	Key capacity metrics

Restriction: You might not be able to view capacity and performance metrics for all of the resources.

For example, you monitor the pools for a SAN Volume Controller and you see that some of the pools are running out of capacity. To notify your colleagues, you export the information about the pools to an HTML file and send it to your colleagues. You want to include information about the input/output performance of the volumes for your storage systems in a report. To create the report, you export the information that is shown in the performance chart to a CSV file.

If you want to extract configuration, capacity, or performance information regularly about your storage resources, you can use the Representational State Transfer (REST) API to automatically export the data that you need to create your reports.

- [Exporting information about capacity shortages in pools](#)
To include information about the capacity of your resources in reports, use the reporting feature in IBM Spectrum® Control.
- [Exporting information about the input/output performance of volumes](#)
To include information about the performance of your resources in reports, use the reporting feature in IBM Spectrum Control.

Related tasks

- [Exporting information to a file](#)

Related information

- [REST API](#)

Exporting information about capacity shortages in pools


To include information about the capacity of your resources in reports, use the reporting feature in IBM Spectrum® Control.

About this task

You check the capacity of pools on your SAN Volume Controller storage system and see that the capacity of some of the pools is nearly depleted. So, you decide to notify your colleagues by including information about the pools in a report.

Procedure

1. Click **Storage > Block Storage Systems**.
2. Double-click the storage system and click **Pools** in the navigation pane.
3. On the Pools page, right-click one of the column headings and click **Zero Capacity**.

Tip: If you want to export information only about the pools that are depleted, click , select **Zero Capacity**, and type **Depleted** in the filter field. When a capacity limit is set, the pool is considered to be **Depleted** when the capacity limit is reached.

4. Click **Actions** and choose one of the options.

Results

You can get information about the capacity of the pools, whether the pools are depleted, and when the pools will become depleted.

Exporting information about the input/output performance of volumes

To include information about the performance of your resources in reports, use the reporting feature in IBM Spectrum® Control.

About this task

You want to include information about the input/output performance of the volumes in a report.

Procedure

1. Click **Storage > Volumes**.
2. Select all of the volumes in the list and click **View Performance**.
By default, the top 10 volumes are ordered by total input/output rate and the reporting period is set to 24 hours.
3. Click **Actions > Export > Export as CSV**.

Results

Information about the performance of the volumes with the highest input/output rates are exported to a CSV file, which you can download and include in a report.

Reporting with Cognos Analytics

Use the Cognos® Analytics reporting tool to view predefined reports and create custom reports about the resources that are monitored by IBM Spectrum® Control.

Charts are automatically generated for most of the predefined reports. Depending on the type of resource, the charts show statistics for space usage, workload activity, bandwidth percentage, and other statistics. You can schedule reports and specify to create the report output in HTML, PDF, and other formats. You can also configure reports to save the report output to your local file system, and to send reports as email attachments.

You must collect information about your environment before you can use reports to view details about the storage resources in it. You can use IBM Spectrum Control monitoring jobs, such as probes and performance monitors, to gather that information.

- [Getting started](#)
You can view predefined reports and create custom reports about IBM Spectrum Control in the optional Cognos Analytics reporting tool. You work with the reports in the Cognos Analytics reporting tool.
- [Scenarios](#)
Use these scenarios to learn how to use IBM Spectrum Control reports to analyze the most active volumes on hypervisors and to monitor the performance of a critical application. You can learn how to investigate a degradation in the performance of a storage pool. You can also learn how to identify the relationships between a storage system and other resources.
- [Configuring](#)
You can configure the optional Cognos Analytics reporting tool according to your requirements.
- [Predefined reports about resource relationships](#)
You can run predefined reports about the end-to-end relationship from a server to back-end storage. You can also run reports about the end-to-end relationship from file systems on servers to volumes on storage pools.
- [Predefined reports about switches and switch ports](#)
You can run predefined reports to analyze and compare the performance of switches and switch ports.
- [Predefined reports about groups](#)
You can run predefined reports to analyze the capacity of groups. A group is a set of logically related volumes, file systems, and shares. For example, a group that represents a business critical application might include the volumes, file systems, and shares that provide storage to the application.
- [Predefined reports about hypervisors](#)
You can run predefined reports to analyze the capacity and performance of hypervisors.
- [Predefined reports about servers and file systems on servers](#)
You can run predefined reports to analyze the capacity and performance of servers, file systems, and volumes.
- [Predefined reports about storage systems and components](#)
You can run predefined reports to view capacity details and performance metrics for storage systems and their components. The components include controllers, modules, nodes, disks, I/O groups, managed disks, storage pools, ports, RAID arrays, and volumes.
- [Predefined reports about storage tiering](#)
Use storage tier reports to help you determine the best storage-tier configuration for your environment.
- [Custom reports about performance](#)
You can use the optional Cognos Analytics reporting tool to create custom performance reports. Performance reports can contain detailed information about the performance of monitored resources, and some of the properties of those resources.
- [Custom reports about capacity and relationships](#)
You can use Cognos Analytics to create custom reports about capacity and relationships. Capacity and relationship reports can contain detailed information about the properties of monitored resources, and the available space and capacity of those resources.
- [Troubleshooting Cognos Analytics reports](#)
Find answers to questions about resolving issues with reports in the optional Cognos Analytics reporting tool.
- [Showing the package version number and build ID](#)
Access information about the packages for custom reports and predefined reports, such as the version number and the build ID.

Getting started

You can view predefined reports and create custom reports about IBM Spectrum® Control in the optional Cognos Analytics reporting tool. You work with the reports in the Cognos Analytics reporting tool.

The following storage systems run IBM Spectrum Virtualize to virtualize their storage: SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, IBM® Storwize® family, FlashSystem 5000, FlashSystem 5100, FlashSystem 7200, FlashSystem 9100, FlashSystem 9200, and FlashSystem V9000. In this documentation, IBM Spectrum Virtualize is used to refer collectively to IBM SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, and IBM Storwize storage systems, and to IBM FlashSystem® devices that run IBM Spectrum Virtualize.

- [Predefined reports](#)
You can use the optional Cognos Analytics reporting tool to view predefined reports about IBM Spectrum Control. Predefined reports provide information about the capacity and performance of your resources.
- [Predefined reports listed by package](#)
The data that the predefined reports use is organized into four packages: Capacity and Relationships, Historical Capacity, Performance, and Storage Tiering. Each report uses data from only one package.
- [Custom reports](#)
You can use the optional Cognos Analytics reporting tool to create custom reports about IBM Spectrum Control. Use custom reports to analyze the capacity and performance of your resources.
- [Viewing predefined reports](#)
Use the Cognos Analytics reporting tool to view predefined reports about IBM Spectrum Control.
- [Creating custom reports](#)
Use the Cognos Analytics reporting tool to view custom reports about IBM Spectrum Control. You can create basic reports with limited formatting, or design custom reports to add advanced formatting features to reports.
- [Viewing and creating reports in the Cognos Analytics reporting tool](#)
You access reports from IBM Spectrum Control, and work with the reports in the Cognos Analytics reporting tool.
- [Searching in your report data](#)
Search through your report data in the Cognos Analytics reporting tool.

Predefined reports

You can use the optional Cognos Analytics reporting tool to view predefined reports about IBM Spectrum® Control. Predefined reports provide information about the capacity and performance of your resources.

Predefined reports

You can view the following types of predefined reports:

Reports about the capacity and relationships of your resources

Reports about the capacity and relationships of servers, hypervisors, network-attached storage systems, storage virtualizers, storage systems, storage resource groups, groups, fabrics, and switches. Use these reports to analyze the capacity and relationships of resources. For example, to show the current space statistics for a pool on a storage system, run the Pools Capacity report. To show the pools that are related to a storage system, run the Storage Resource Relationships Summary report.

Reports about performance

Reports about the performance of your monitored resources. You can view reports about storage systems, components of storage systems, switches, and switch ports. Use these reports to analyze the performance of resources. For example, to find out which are the most active volumes on a storage system over a specified period, run the Most Active Volumes report. To view details about the performance of a volume over a specified period, run a report such as the Performance of One Volume report.

Restrictions:

- Performance data is collected at intervals. An interval represents the number of minutes over which samples of performance data are averaged. Data that is collected at certain intervals is automatically consolidated, or rolled up, to higher intervals. For example, data collected at 1-minute intervals is consolidated into 5-minute data. When you create a performance report in the Cognos Analytics reporting tool, the smallest interval for the sample data is 5 minutes. The data collected at 1-minute intervals is not available in the Cognos Analytics reporting tool to avoid performance problems.
- Performance reports are not available for the following storage systems:
 - FlashSystem 900
 - IBM Spectrum Scale
 - Dell EMC Unity
 - Hitachi VSP
 - NetApp storage systems running ONTAP 9
 - Pure FlashArray//M and FlashArray//X storage systems

To view performance metrics for those storage systems, use the IBM Spectrum Control GUI. For more information, see [Monitoring the performance of resources](#).

Reports about historical capacity

Reports that include the historical information about the capacity of your storage systems, pools, and volumes. Use these reports to determine whether more storage resources are required. For example, to show how space is used over a specified period on volumes on storage systems, run the Volumes Historical Capacity report. You can also use these reports to analyze future storage requirements for your environment.

Reports about storage tiering

Reports that include information about the workload activity of these resources: managed disk (MDisk) groups, virtual disks (VDisks) in MDisk groups, VDisks in storage resource groups, and VDisks in storage virtualizers. Use these reports to determine the best storage-tier configuration for your environment.

Charts are automatically generated for most of the predefined reports. Depending on the type of resource, the charts show statistics for space usage, workload activity, bandwidth percentage, and other statistics.

Tip: When you view a predefined report, click Help to find out more about the report.

Related reference

- [Predefined reports listed by package](#)

Predefined reports listed by package

The data that the predefined reports use is organized into four packages: Capacity and Relationships, Historical Capacity, Performance, and Storage Tiering. Each report uses data from only one package.

Tip: When you view a predefined report, click Help to find out more about the report.

Capacity and Relationships package

You can view the following predefined reports about capacity and relationships:

Availability of Switch Ports

Shows a chart of the 20 switches with the most ports and connected ports, and a table with the properties of the switches.

Capacity of One Group

Shows charts and tables of the capacity, available space, and used space statistics for each type of resource in a group. The report shows separate charts and tables for block-level and file-level resources.

Disks Capacity

Shows a chart of the disk classes with the most allocated space, and a table with space statistics for all disks by class. Use the report to review space allocation in your disks.

File System to Volume Relationships

Shows the end-to-end relationship from file systems on servers to volumes on storage pools. The report shows the usage of space on file systems over a specified period. Use the report to analyze changes in your file system requirements.

File Systems Capacity

Shows the available and used capacity of file systems, which are grouped by server. You can use this report to identify file systems that might be running out of space.

Groups Capacity

Shows a chart of the 20 groups with the most capacity, available space, or used space, and a table with details of all groups. Use the report to review the space allocation in your groups.

Hypervisor Data Stores Capacity

Shows a chart of the 20 hypervisor data stores with the most allocated space. This report also shows a table with space statistics for all hypervisor data stores. Use the report to review space allocation in your hypervisor data stores.

Hypervisor Disks Capacity

Shows a chart of the 20 hypervisor disks with the most allocated space, and a table with space statistics for all hypervisor disks. Use the report to review space allocation in your hypervisor disks.

Hypervisors Capacity

Shows a chart of the 20 hypervisors with the most allocated space, and a table with space statistics for all hypervisors. Use the report to review space allocation in your hypervisors.

Managed Disks Capacity

Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

Pools Capacity

Shows a chart of the 20 pools with the most allocated space, and a table with space statistics for all pools. Use the report to review space allocation in your pools.

Server Disks Capacity

Shows the available and used disk space for all disks, which are grouped by server. You can use the report to identify disks that have no allocated space.

Servers Capacity

Shows a summary of the available and used disk space, and the available and used file systems, for servers. Use the report to identify servers that are running out of space.

Storage Resource Relationships Summary

Shows the end-to-end relationship from a server to back-end storage. The report shows servers, hypervisors, NAS, storage systems, storage virtualizers, pools, and volumes.

Storage Resource Relationships Summary (Configurable)

Shows the end-to-end relationship from a server to back-end storage. You can configure which resources to show in the report.

Storage Systems Capacity

Shows a chart of the 20 storage systems with the most used space, and a table with details of all storage systems. Use the report to review the space allocation in your storage systems.

Volumes Capacity

Shows a chart of the 20 volumes with the most used space, and a table with details of all volumes. Use the report to review the space allocation in your volumes.

Historical Capacity package

You can view the following predefined reports about historical capacity:

Pools Historical Capacity

Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

Storage Systems Historical Capacity

Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

Volumes Historical Capacity

Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

Performance package

Performance data is collected at intervals. An interval represents the number of minutes over which samples of performance data are averaged. Data that is collected at certain intervals is automatically consolidated, or rolled up, to higher intervals. For example, data collected at 1-minute intervals is consolidated into 5-minute data. When you create a performance report in the Cognos® Analytics reporting tool, the smallest interval for the sample data is 5 minutes. The data collected at 1-minute intervals is not available in the Cognos Analytics reporting tool to avoid performance problems.

You can view the following predefined reports about performance:

- Compare Performance of Multiple Controllers and Modules
 - Shows up to four performance metrics for multiple controllers or modules over time. Use the report to compare the performance metrics for controllers or modules that are on storage systems.
- Compare Performance of Multiple Disks
 - Shows up to four performance metrics for multiple local disks over time. Use the report to compare the performance metrics for local disks that are on storage systems.
- Compare Performance of Multiple IO Groups
 - Shows up to four performance metrics for multiple I/O groups over time. Use the report to compare the performance metrics for I/O groups that are on storage systems.
- Compare Performance of Multiple Managed Disks
 - Shows up to four performance metrics for multiple managed disks over time. Use the report to compare the performance metrics for managed disks that are on storage systems.
- Compare Performance of Multiple Nodes
 - Shows up to four performance metrics for multiple nodes over time. Use the report to compare the performance metrics for nodes that are on storage systems.
- Compare Performance of Multiple Pools
 - Shows up to four performance metrics for multiple storage pools over time. Use the report to compare the performance metrics for storage pools that are on storage systems.
- Compare Performance of Multiple Ports
 - Shows up to four performance metrics for multiple ports over time. Use the report to compare the performance metrics for ports that are on storage systems.
- Compare Performance of Multiple Storage Systems
 - Shows up to four performance metrics for multiple storage systems over time. Use the report to compare the performance metrics for storage systems.
- Compare Performance of Multiple RAID Arrays
 - Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.
- Compare Performance of Multiple Switch Ports
 - Shows up to four performance metrics for multiple switch ports over time. Use the report to compare the performance metrics for switch ports.
- Compare Performance of Multiple Switches
 - Shows up to four performance metrics for multiple switches over time. Use the report to compare the performance metrics for switches.
- Compare Performance of Multiple Volumes
 - Shows up to four performance metrics for multiple volumes over time. Use the report to compare the performance metrics for volumes that are on storage systems.
- Compare Performance of One Pool over Time Ranges
 - Shows one performance metric on one pool over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.
- Compare Performance of One Storage System over Time Ranges
 - Shows one performance metric on one storage system over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.
- Compare Performance of One Switch over Time Ranges
 - Shows one performance metric on one switch over two time periods. For example, you can use the report to compare the I/O rate for this week and last week.
- Compare Performance of One Switch Port over Time Ranges
 - Shows one performance metric on one switch port over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.
- Most Active Controllers or Modules
 - Shows a chart of the 20 controllers or modules that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of controllers or modules.
- Most Active Disks
 - Shows a chart of the 20 local disks that are most active, and details for all local disks, for a time period that you specify. Use the report to analyze the performance of local disks.
- Most Active Host Connections
 - Shows a chart of the 20 host connections that have the most active volumes. The most active volumes have the greatest aggregate load.
- Most Active Hypervisors
 - Shows a chart of the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load.
- Most Active IO Groups
 - Shows a chart of the 20 I/O groups that are most active, and details for all I/O groups, for a time period that you specify. Use the report to analyze the performance of I/O groups.
- Most Active Managed Disks
 - Shows a chart of the 20 managed disks that are most active, and details for all managed disks, for a time period that you specify. Use the report to analyze the performance of managed disks.
- Most Active Nodes
 - Shows a chart of the 20 nodes that are most active, and details for all nodes, for a time period that you specify. Use the report to analyze the performance of nodes.
- Most Active Pools
 - Shows a chart of the 20 storage pools that are most active, and details for all pools, for a time period that you specify. Use the report to analyze the performance of pools.
- Most Active Ports
 - Shows a chart of the 20 ports that are most active, and details for all ports, for a time period that you specify. Use the report to analyze the performance of ports.
- Most Active RAID Arrays
 - Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.
- Most Active Servers
 - Shows a chart of the 20 servers that have the most active volumes. The most active volumes have the greatest aggregate load.
- Most Active Storage Systems
 - Shows a chart of the 20 storage systems that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of storage systems.
- Most Active Switch Ports
 - Shows a chart of the 20 switches that are most active, and details for all switches, for a time period that you specify. Use the report to analyze the performance of switches.
- Most Active Volumes

Shows a chart of the 20 volumes that are most active, and details for all volumes, for a time period that you specify. Use the report to analyze the performance of volumes.

Performance Data Export

Exports the performance metrics for one or more resources on a storage system to a spreadsheet file in Microsoft Excel file format. Use the report if you want to analyze performance data in Excel or to send performance data to IBM®.

Performance of One Controller or Module

Shows four charts, and a table of performance metrics for a controller or module over a time period. Use the report to analyze the performance of a controller or module.

Performance of One Disk

Shows four charts, and a table of performance metrics for a local disk over a time period. Use the report to analyze the performance of a local disk.

Performance of One IO Group

Shows four charts, and a table of performance metrics for an I/O group over a time period. Use the report to analyze the performance of an I/O group.

Performance of One Managed Disk

Shows four charts, and a table of performance metrics for a managed disk over a time period. Use the report to analyze the performance of a managed disk.

Performance of One Node

Shows four charts, and a table of performance metrics for a node over a time period. Use the report to analyze the performance of a node.

Performance of One Pool

Shows four charts, and a table of performance metrics for a storage pool over a time period that you specify. Use the report to analyze the performance of a storage pool.

Performance of One Port

Shows four charts, and a table of performance metrics for a port over a time period that you specify. Use the report to analyze the performance of a port.

Performance of One RAID Array

Shows four charts, and a table of performance metrics for a RAID array over a time period. Use the report to analyze the performance of a RAID array.

Performance of One Storage System

Shows four charts, and a table of performance metrics for a storage system over a time period that you specify. Use the report to analyze the performance of a storage system.

Performance of One Switch

Shows three charts, and a table of performance metrics for a switch over a time period that you specify. Use the report to analyze the performance of a switch.

Performance of One Switch Port

Shows four charts, and a table of performance metrics for a switch port. Use this report to view multiple performance metric types on the same chart. For example, you can view I/O rate and response time metrics on the same chart.

Performance of One Volume

Shows four charts, and a table of performance metrics for a volume over a time period that you specify. Use the report to analyze the performance of a volume.

Performance of Volumes by Host Connection

Shows the performance metrics for volumes on a host connection. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

Performance of Volumes by Hypervisor

Shows the performance metrics for volumes on a hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.

Performance of Volumes by Server

Shows the performance metrics for volumes on a server. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

Summarized Performance of Volumes by Host Connection

Shows summarized performance metrics for volumes on a host connection. Use the report to identify the aggregate load and the average response time of a particular server.

Summarized Performance of Volumes by Hypervisor

Shows summarized performance metrics for volumes on a hypervisor. Use the report to identify the aggregate load and the average response time of a particular hypervisor.

Summarized Performance of Volumes by Server

Shows summarized performance metrics for volumes on a server. Use the report to identify the aggregate load and the average response time of a particular server.

Storage Tiering package

You can view the following predefined reports about storage tiering:

MDisk Group - VDisk Workload Activity Details

Shows the workload activity of virtual disks (VDisks) in a managed disk (MDisk) group. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

MDisk Group Details

Shows the workload activity of a managed disk (MDisk) group. Use the report to analyze storage services for an MDisk group.

MDisk Groups - VDisk Workload Activity

Use this report to monitor the performance of virtual disks (VDisks) in managed disk (MDisk) groups. The workload activity of each MDisk group is determined by calculating the average and the maximum peak utilization of the VDisks in the MDisk group.

MDisk Groups - Workload Activity

Shows the workload activity of managed disk (MDisk) groups. You can use the information that the report provides to investigate which MDisk groups are being underused or overused.

Storage Resource Group - VDisk Workload Activity

Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

Storage Resource Group - VDisk Workload Activity Details

Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

Storage Virtualizer - VDisk Workload Activity

Shows the workload activity of virtual disks (VDisks) in a storage virtualizer. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

VDisk Details

Shows the workload activity of a virtual disk (VDisk) in a managed disk group. Use the report to analyze the performance of a VDisk in a managed disk (MDisk) group.

Custom reports

You can use the optional Cognos Analytics reporting tool to create custom reports about IBM Spectrum® Control. Use custom reports to analyze the capacity and performance of your resources.

Custom reports

Unlike predefined reports, you can select the information that you want to include in custom reports. Custom reports can contain detailed information about the relationships between monitored resources, the properties of monitored resources, and detailed information about the performance of monitored resources.

To create a report, browse the list of monitored resources in your storage environment, and drag items that are associated with the resources into a report. For example, drag items that show capacity data for storage systems into a report.

If you create custom reports, you can also do the following tasks:

- Copy a predefined report and modify the copy to create a custom report.
- Create custom charts. You can create charts with data that is in the combinations that you want to see. Use these custom charts to troubleshoot and to report on the status of resources.
- Create reports with multiple charts. You can create multiple charts with different scales, and display those charts side-by-side. You can use these charts to compare performance metrics. For example, you can review read I/O rates and the number of bytes transferred in the same report.
- Create a custom folder structure to organize your custom reports.
- Run a report from a URL. You can include options for the report in the URL. For options that are not included in the URL, you can specify default values, or prompt the user to specify a value.
- Share a custom report with your organization, or keep it for only your own use. You can also distribute different sections of a report to different teams.

Tip: Custom reports are stored in the Content Manager (cm) database. Back up this database regularly to prevent the loss of your custom reports if a problem occurs with the cm database.

Advanced formatting for custom reports

You can create reports with basic formatting, or design more complex reports with advanced formatting.

Access related reports

You can configure custom reports so that you can click a resource in the report output to access other related reports. If you identify a problem in a report, you can open another report from within that report that might provide specific information about the problem.

For example, the Most Active Servers predefined report shows data for servers. Click the name of a server in the report to view the Summarized Performance of Volumes by Server report for the server. This report shows the aggregated data for all of the volumes on the server. Click the name of a volume in the report to view the Performance of Volumes by Server report for the volume.

Use GUI controls in reports

You can configure custom reports to include GUI controls to enable users to select options for the reports. You can specify default values for these options. In addition, you can create copies of reports with different default values.

The GUI controls that you can add to your reports enable users to select dates, type values, and select from lists. You can also create GUI controls whose values depend on what is selected in another GUI control. You can include GUI controls on the report output in HTML format so that users can change options and run the report again. For example, you can include a drop-down list in the HTML report output so that users can change the sort order of the data.

Schedule reports

You can schedule when to run the report. You can specify the following formats for the report: CSV, Excel, HTML, PDF, and XML. You can specify to email the report output and to save the report output to a location on your file system or on a web server. For example, you might require a monthly report that shows the usage of space on storage systems. Management staff might use this report to analyze changes in your storage system requirements. You can create this report and schedule it to run every month. You can specify that the report output is emailed to management staff, or saved to a location from which management staff can access it.

Viewing predefined reports

Use the Cognos® Analytics reporting tool to view predefined reports about IBM Spectrum® Control.

Before you begin

If you use Internet Explorer 10 or 11, you must set the browser to compatibility mode. For information about compatibility mode, see <http://windows.microsoft.com/en-us/internet-explorer/use-compatibility-view>.

Procedure

1. Go to the URL for your Cognos Analytics server. The format of the URL is similar to this URL: `http://myhostname:9300/bi`
2. In the Welcome portal, click Team Content.
3. Click IBM Spectrum Control Predefined Reports.
4. Navigate to the report that you want to see, and then click the name of the report.

To view a description of a report or folder, right-click the report or folder, then click Properties. The description is in the General tab.

Tip: When you run some predefined reports, you can select the configuration of the storage systems. You can select to run the report on systems that are configured as storage systems or as storage virtualizers. If you select one of these configurations, you must also select one or more storage systems from the list of available storage systems. If you do not specify storage systems, the report shows data for both configurations.

Creating custom reports

Use the Cognos® Analytics reporting tool to view custom reports about IBM Spectrum® Control. You can create basic reports with limited formatting, or design custom reports to add advanced formatting features to reports.

About this task

Restriction:



You cannot create custom reports from the following packages:

- Historical Capacity
- Storage Tiering

Note: If you use Microsoft Internet Explorer 10 or 11, you must set the browser to compatibility mode. For information about compatibility mode, go to the following link:

➔ <http://windows.microsoft.com/en-us/internet-explorer/use-compatibility-view>

Procedure

1. Go to the URL for your Cognos Analytics server. The format of the URL is similar to this URL: `http://myhostname:9300/bi`
2. Depending on the type of report that you want to see, open the appropriate report tool:
 - To create a basic report with limited formatting about the capacity and relationships of resources, complete the following steps:
 - Click New  in the Welcome portal.
 - Click Other, then click Query Studio.
 - Click IBM Spectrum Control Packages, then click the Capacity and Relationships package.
 - To create a basic report with limited formatting about the performance of resources, complete the following steps:
 - Click New  in the Welcome portal.
 - Click Other, then click Query Studio.
 - Click IBM Spectrum Control Packages, then click the Performance package.
3. Optional: To create a report with advanced formatting features, complete the following steps:
 - a. Click Team Content in the Welcome portal.
 - b. Click IBM Spectrum Control Packages.
 - c. Right-click Capacity and Relationships or Performance, then click Create report.
 - d. Click a template for the report.

Related tasks

- [Creating custom performance reports](#)
- [Creating custom capacity and relationship reports](#)
- [Adding filters to reports](#)


Viewing and creating reports in the Cognos Analytics reporting tool

You access reports from IBM Spectrum® Control, and work with the reports in the Cognos® Analytics reporting tool.

About this task

When you view and create reports, you can navigate from predefined reports to custom reports and from custom reports to predefined reports within the Cognos Analytics reporting tool.

Procedure

1. In the Cognos Analytics reporting tool, click Team Content in the Welcome portal.
2. Depending on the type of report that you want to work with, complete one of the following actions:
 - To view predefined reports, click IBM Spectrum Control Predefined Reports.
 - To create a basic custom report with limited formatting, complete the following steps:
 - Click the New icon  in the Welcome portal.
 - Click Other, then click Query Studio.
 - Click IBM Spectrum Control Packages, then click the Capacity and Relationships or Performance package.
 - To create a report and apply advanced formatting features to the report, complete the following steps:
 - Click IBM Spectrum Control Packages.
 - Right-click the Capacity and Relationships or Performance package, then click Create report.
 - Click a template for the report.

Searching in your report data

Search through your report data in the Cognos® Analytics reporting tool.

Procedure

1. In the Welcome portal, click Search.
2. Type the search terms in the Search field, then press Enter.
Tip: You can filter and save searches.
3. Optional: Alternatively, you can search the online help for properties or metrics.
 - a. In the Cognos Analytics reporting tool, click any predefined report.
 - b. Click Help, which is in the upper-right corner of the page.
 - c. In the Search IBM Documentation field, type the name of the property or metric in quotation marks followed by the word `report`. Press Enter to search for the term.
For example, type one of the following:
 - "Storage Pool Name" report
 - "Port Receive I/O Rate" reportThe list of help topics that contain the property or metric is displayed in the search results.

Scenarios

Use these scenarios to learn how to use IBM Spectrum Control reports to analyze the most active volumes on hypervisors and to monitor the performance of a critical application. You can learn how to investigate a degradation in the performance of a storage pool. You can also learn how to identify the relationships between a storage system and other resources.

- [Analyzing the volumes on the most active hypervisors](#)
In this scenario, you use the predefined reports for hypervisors to analyze the volumes on hypervisors. You can view the most active volumes. Then, you can view summary information for a hypervisor and details for individual volumes. You can also use this scenario to analyze the volumes on servers, agentless servers, and cluster resource groups.
- [Monitoring daily the performance of volumes](#)
In this scenario, you create a report view from a predefined report to run a daily report about the performance of volumes. You can also use this scenario to monitor any other resources and performance metrics.
- [Investigating a degradation in the performance of a storage pool](#)
You can use the Compare Performance of One Pool over Time Ranges report to analyze a change in the performance of a storage pool. You can also use this scenario to compare the performance metrics of storage systems, switches, and switch ports over two periods of time.
- [Identifying the relationships between a storage system and other resources](#)
In this scenario, you use a predefined relationships report to identify the relationships between a storage system and other servers or hypervisors. You can also use this scenario to identify the relationships between storage systems and NAS, pools, and volumes. Use this scenario to identify the relationships between storage virtualizers and servers, hypervisors, NAS, pools, and volumes.

Analyzing the volumes on the most active hypervisors

In this scenario, you use the predefined reports for hypervisors to analyze the volumes on hypervisors. You can view the most active volumes. Then, you can view summary information for a hypervisor and details for individual volumes. You can also use this scenario to analyze the volumes on servers, agentless servers, and cluster resource groups.

About this task

You have 10-20 volumes on the hypervisors in your storage environment. You receive reports that the performance of some volumes is slow. You want to analyze the performance of these volumes as a group, and the performance of individual volumes.

Procedure

1. Run the Most Active Hypervisors report.
Use the chart in the report output to see the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load. You see that the volumes on one hypervisor are considerably more active than the volumes on other hypervisors. You want to investigate that hypervisor.
2. To analyze the activity on this hypervisor, click the name of the hypervisor in the report table.
The Summarized Performance of Volumes by Hypervisor report runs for the hypervisor. Use the report to see summarized performance metrics for the volumes on the hypervisor. You can also use the report to see the aggregate load and the average response time of the volumes on the hypervisor.
In the charts and the table on the report, you see that the I/O response times for the hypervisor are slower than normal.
3. To investigate the individual volumes on the hypervisor, click the name of the hypervisor in the report table.
The Performance of Volumes by Hypervisor report runs for the hypervisor. Use the report to see the performance metrics for individual volumes on the hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.
In the charts and the table on the report, you see that the I/O response times for one volume are slow.

What to do next

You identified the volume with the slow I/O response times. You can now copy data to the other volumes to more evenly distribute the load of the volumes.

Related reference

- [Most Active Hypervisors report](#)
- [Summarized Performance of Volumes by Hypervisor report](#)
- [Performance of Volumes by Hypervisor report](#)

Monitoring daily the performance of volumes

In this scenario, you create a report view from a predefined report to run a daily report about the performance of volumes. You can also use this scenario to monitor any other resources and performance metrics.

About this task

You have applications that are critical to your business. The applications access volumes in your storage environment. You want to monitor daily the performance of the volumes to confirm that they are functioning correctly. To monitor the volumes, you must create a report view to run a daily report about the performance of volumes.

Procedure

1. Go to the URL for your Cognos® Analytics server. The format of the URL is similar to this URL: `http://myhostname:9300/bi`
2. In the Welcome portal, click Team Content.
3. Click IBM Spectrum Control Predefined Reports.
4. Click Storage Systems > Volumes.
5. Right-click the Compare Performance of Multiple Volumes report, then click Create report view.
6. Type a name for the report view, then click Save.
7. Right-click the new report view, then click Properties.
8. Click the Schedule tab.
9. Click New to create new schedule.
10. Schedule the report view to run daily.
11. In the Options area, click Format, then select HTML and PDF formats.
12. In the Prompts area, click Set values, and then click Set.
13. Specify the storage systems and the volumes that you want to monitor, and then click Next.
14. Select the following performance metrics:
 - Total Overall I/O Rate (ops/s)
 - Overall Response Time (ms/op)
 - Total Data Rate (MiB/s)
 - Total Overall Cache Hit Percentage
15. Specify an interval and a reporting period for the report view, and then click Finish.
16. Click Create to save the schedule, format, and prompt value settings.

What to do next

View the output of this report every day to confirm that these volumes are functioning correctly. Click the report view to view the latest output for the report. If the report output is not displayed, make sure that the Default action for report view in the properties of the report view is set to View most recent report.

Related reference

- [Compare Performance of Multiple Volumes report](#)

Investigating a degradation in the performance of a storage pool

You can use the Compare Performance of One Pool over Time Ranges report to analyze a change in the performance of a storage pool. You can also use this scenario to compare the performance metrics of storage systems, switches, and switch ports over two periods of time.

About this task

In the previous three days, the performance of a storage pool has significantly degraded. You want to investigate the reasons for the degradation. To do this, use the Compare Performance of One Pool over Time Ranges report to compare the recent poor performance of the pool with a period when the performance was normal.

Procedure

1. Go to the URL for your Cognos® Analytics server. The format of the URL is similar to this URL: `http://myhostname:9300/bi`
2. In the Welcome portal, click Team Content.
3. Click IBM Spectrum Control Predefined Reports.
4. Click Storage Systems, and then click Pools.
5. Click Compare Performance of One Pool over Time Ranges.
6. Specify the storage system and the pool whose performance has degraded.
7. Specify the performance metric that you want to analyze, for example Overall Response Time (ms/op).
8. Specify an interval.
9. Specify two custom date ranges for the two reporting periods. Specify the previous three days as one reporting period. For the other reporting period, specify a period when the performance of the pool was normal.
10. Click Finish.

What to do next

To investigate other possible reasons for the degradation, run the report again, and select another metric.

Related reference

- [Compare Performance of One Pool over Time Ranges report](#)

Identifying the relationships between a storage system and other resources

In this scenario, you use a predefined relationships report to identify the relationships between a storage system and other servers or hypervisors. You can also use this scenario to identify the relationships between storage systems and NAS, pools, and volumes. Use this scenario to identify the relationships between storage virtualizers and servers, hypervisors, NAS, pools, and volumes.

About this task

You want to take a storage system offline and migrate the volumes to another storage system. To prepare for the migration, you first identify which servers or hypervisors are affected when the storage system goes offline.

Procedure

1. Go to the URL for your Cognos® Analytics server. The format of the URL is similar to this URL: `http://myhostname:9300/bi`
2. In the Welcome portal, click Team Content.
3. Click IBM Spectrum Control Predefined Reports.
4. Click Storage Resource Relationships Summary.
5. From the Storage systems list, select the storage system that you plan to take offline.
6. Click Finish.

What to do next

View the output of the report to see which servers or hypervisors are affected when the storage system goes offline. You can now migrate the volumes.

Related reference

- [Storage Resource Relationships Summary report](#)

Configuring

You can configure the optional Cognos Analytics reporting tool according to your requirements.

- [Setting options for reports](#)
When you schedule or run a report, you can specify the output format and how the report is delivered.
- [Scheduling a report to run with predefined options](#)
You can schedule a report to run on different systems and with different options that you define. To monitor the performance of different sets of volumes, you can schedule a Compare Performance of Multiple Volumes report for each of the sets.
- [Creating a copy of a predefined report](#)
Create a copy of a predefined report if you want to modify report. You might want to modify the queries that the report uses, or to add or remove a column. Changes to the predefined report do not affect copies of the report.
- [Starting the Cognos Configuration GUI](#)
You can use the IBM® Cognos Configuration GUI to configure Cognos Analytics reporting tool components. For example, you can configure an archive location for the output of reports, and configure a connection between the Cognos Analytics reporting tool and an SMTP server.
- [Sending reports by email](#)
To send reports as attachments in email messages, configure a connection between the Cognos Analytics reporting tool and an SMTP server.
- [Specifying locations for saving reports](#)
You can configure reports to save the report output to your local file system. You specify a root directory in the file system in which the output of reports can be saved in the IBM Cognos Configuration GUI. You specify the locations in which you want to save the output of reports. Then, you can configure individual reports to save the report output to a location in the file system. For example, you can schedule a report and configure the report output to be saved to a location like C:\marketing_team\sales or /usr/marketing_team/sales.
- [Setting up multiple connections to the database repository](#)
In Cognos Analytics, you can add multiple connections to the IBM Spectrum Control data source so that you can run reports against multiple IBM Spectrum Control databases (TPCDBs).
- [Customizing the logo and title of reports](#)
You can change the logo that is displayed in all IBM Spectrum Control reports and the title of individual reports. Before you modify predefined reports, create backup copies of them.

Setting options for reports

When you schedule or run a report, you can specify the output format and how the report is delivered.

About this task

You can select output formats for the report. For example, you can select PDF, HTML, and other formats. If you choose PDF as the format for your report, you can set options such as the orientation and paper size for reports.

You can also specify to print the report, to send the report by email, and other delivery options.

Procedure

1. Navigate to a report in the Cognos® Analytics reporting tool.
2. Depending on whether you want to schedule a report or run a report, do one of the following steps:
 - To schedule a report and select format and delivery options, right-click the report that you want to schedule, then click Properties. Click the Schedule tab, then click New.
 - To run a report and select format and delivery options, right-click the report, then click Run as. Click Run in background, then click Advanced to select format and delivery options.

Related information

- [Schedule an Entry](#)
- [Report Formats](#)
- [Distributing Reports](#)
- [Advanced Report Options](#)

Scheduling a report to run with predefined options

You can schedule a report to run on different systems and with different options that you define. To monitor the performance of different sets of volumes, you can schedule a Compare Performance of Multiple Volumes report for each of the sets.

About this task

You can specify schedules and output formats for the reports, and different ways to share the reports.

To schedule a report to run with parameters that you specify, you first create a report view from the report. Then, you schedule the report and set the parameters and options in the report view.

If the report from which a report view was created is updated, the report view is also updated.

Procedure

1. Go to the URL for your Cognos® Analytics server. The format of the URL is similar to this URL: **http://myhostname:9300/bi**
2. In the Welcome portal, click Team Content.
3. Click IBM Spectrum Control Predefined Reports.
4. Right-click the report, then click Create report view.
5. Type a name for the report view, then click Save.
6. To set predefined options for the report view, right-click the report view, complete one or more of the following steps:
 - To set a schedule for the report view, right-click the report view, and then click the Properties tab. Add a new schedule for the report view in the Schedule tab.
 - To set formats for the report view, right-click the report view, and then click the Properties tab. Set the format for the report view in the Report tab. Click Report options to set the format options.
 - You can specify parameters such as the resources, the interval, and the reporting period for the report view. To set parameters for the report view, right-click the report view, and then click the Properties tab. In the Report tab, click Set values. Click Set on the Current values panel. Select the values for the parameters, then click Finish.

Related tasks

- [Monitoring daily the performance of volumes](#)

Related information

- [Report Views](#)

Creating a copy of a predefined report

Create a copy of a predefined report if you want to modify report. You might want to modify the queries that the report uses, or to add or remove a column. Changes to the predefined report do not affect copies of the report.

About this task

When you create a copy of a report, the links to related reports in the report output do not work correctly.

If the predefined report is changed during an upgrade, the copy of the report is not affected. If you modify a predefined report instead of creating a copy and then you upgrade your reports software, your modifications might be overwritten.

Procedure

1. In the Welcome portal, click Team Content.
2. Click IBM Spectrum Control Predefined Reports.
3. Navigate to the report that you want to copy, right-click the report, then click Copy or move.
4. Navigate to the location where you want to copy the report, then click Copy to.

Related information

- [Copy a Report](#)

Starting the Cognos Configuration GUI

You can use the IBM® Cognos® Configuration GUI to configure Cognos Analytics reporting tool components. For example, you can configure an archive location for the output of reports, and configure a connection between the Cognos Analytics reporting tool and an SMTP server.

Procedure

To start Cognos Configuration GUI, follow the instructions for your operating system:

Windows

- Click Start, IBM Cognos Configuration.
- Alternatively, run the `cogconfig.bat` file in the `cognos_analytics_install_directory\bin64` directory.

AIX® or Linux®

Log on to a graphical session, and then run the following commands:

```
cd cognos_analytics_install_directory/bin64/  
./home/db2inst1/sqllib/db2profile  
./cogconfig.sh
```

Note the space between the `.` and the `/` in the second command.

Sending reports by email

To send reports as attachments in email messages, configure a connection between the Cognos® Analytics reporting tool and an SMTP server.

About this task

You use the IBM® Cognos Configuration GUI to configure a connection between the Cognos Analytics reporting tool and an SMTP server. When you configure the connection to the SMTP server, you can send the custom and predefined reports that you create as email attachments.

Procedure

1. Open the IBM Cognos Configuration GUI.
2. On the Explorer pane, click Data Access.
3. Click Notification, and then enter values in the following fields:
 - SMTP Mail Server
 - Enter the host name and port number of the mail server.
 - Account and password
 - Leave the field blank if authentication is not required to log on to the mail server. If authentication is required to log on to the mail server, click the Edit icon.
 - Enter a valid user ID and password, and then click OK.
 - Default sender
 - Enter the email address of the sender.
4. To ensure that the connection to the mail server is correctly configured, right-click Notification, and then click Test.
5. Click Save.
6. Click File, Exit.
7. To restart the IBM Cognos server, click Yes in the IBM Cognos Configuration message.

Related tasks

- [Starting the Cognos Configuration GUI](#)

Specifying locations for saving reports

You can configure reports to save the report output to your local file system. You specify a root directory in the file system in which the output of reports can be saved in the IBM® Cognos® Configuration GUI. You specify the locations in which you want to save the output of reports. Then, you can configure individual reports to save the

report output to a location in the file system. For example, you can schedule a report and configure the report output to be saved to a location like C:\marketing_team\sales or /usr/marketing_team/sales.

- [Specifying a root directory for saving reports](#)
You can specify a root directory in the file system in which the output of reports can be saved. To do this, you configure an archive location in IBM Cognos Configuration.
- [Specifying file system locations for reports](#)
You can specify the locations in which you want to save the output of reports in the Cognos Analytics reporting tool.
- [Specifying a report to save the report output to a file system](#)
You can configure a report to save the report output to a location in a file system. You can also specify the file name for the report output, and whether to append a date and time to the file name.

Specifying a root directory for saving reports

You can specify a root directory in the file system in which the output of reports can be saved. To do this, you configure an archive location in IBM® Cognos® Configuration.

Procedure

1. Create a directory in your file system to use as the root directory for the output of reports.
Tip: Ensure that the directory is accessible to users and separate from the Cognos Analytics installation directory. For example, in a distributed installation on Microsoft Windows, you can use an archive folder such as \\server_name\directory.
2. Open the IBM Cognos Configuration GUI.
3. Click Actions > Edit Global Configuration.
4. Click the General tab.
5. In the Archive Location File System Root field, enter a URI using the format file:///directory where *directory* is the directory you created in step 1.
The file:/// part of the URI is required. You can use Windows Universal Naming Convention (UNC) names, such as \\server_name\directory. If you use a Windows UNC name, format the URI as follows:

file:///\\server_name\directory
Tip: Don't use a mapped drive when you run Cognos Analytics as a Microsoft Windows service.
6. Click Test to confirm that the correct location will be used.
7. Click OK.
8. Click File > Save.
9. Click Actions > Restart or Actions > Start to restart the Cognos Analytics server.

Related tasks

- [Starting the Cognos Configuration GUI](#)

Specifying file system locations for reports

You can specify the locations in which you want to save the output of reports in the Cognos® Analytics reporting tool.


Before you begin

You must configure the root directory of the locations that you specify in the Cognos Analytics reporting tool. You must stop and restart the Cognos Analytics server after you specify the locations.

About this task

Specify the locations for report output if you want to share reports with an external application, such as a web site. The reports are saved to this location every time they are updated so that current content is always available. You can also save reports on a local network for users who do not have access to Cognos Analytics software. You can specify multiple locations for the report output.

Procedure

1. Go to the URL for your Cognos Analytics server. The format of the URL is similar to this URL: `http://myhostname:9300/bi`
2. Click Manage in the Welcome portal, then click Administration console.
3. Click the Configuration tab.
4. Click Dispatchers and Services.
5. Click the Define File System Locations icon .
6. On the Define file system locations page, click New.
7. Specify a name and the directory for the location on the file system, and then click Finish.
For example, enter `sales`.
The directory is created the first time that the output of a report for the directory is created.
Repeat this step to configure more file system locations.

Results

The location that you specified is appended to the file system root location that is specified in the IBM® Cognos Configuration GUI. In the IBM Cognos Configuration GUI, the file system root location is referred to as the Archive Location File System Root.

When users select Save report as an external file as the report delivery method when they run or schedule a report, the report output files are saved to this location each time the report runs.

Related tasks

- [Starting the Cognos Configuration GUI](#)

Specifying a report to save the report output to a file system

You can configure a report to save the report output to a location in a file system. You can also specify the file name for the report output, and whether to append a date and time to the file name.

Before you begin

You must configure the locations in which you want to save the output of reports in the Cognos® Analytics reporting tool.

Procedure

1. In the Cognos Analytics reporting tool, right-click the report, then click Properties.
2. Click the Schedule tab.
3. Click New to create a new schedule.
4. In the Options area, click Delivery, then click Save report as an external file.
5. Optional: Specify the name to use for the report, the location of the report output on the file system, and what to do if there is a file name conflict.
6. Click Done.
7. Click Create to save the schedule for the report.

Setting up multiple connections to the database repository

In Cognos® Analytics, you can add multiple connections to the IBM Spectrum® Control data source so that you can run reports against multiple IBM Spectrum Control databases (TPCDBs).

Before you begin


If your IBM Spectrum Control database is on a different computer than the computer where you installed Cognos Analytics, you must catalog the node and database for the remote IBM Spectrum Control database before you create the Cognos Analytics data source for the IBM Spectrum Control database.

About this task

If you create multiple connections, you can run reports for multiple IBM Spectrum Control servers from one Cognos Analytics server. You can also develop and test reports on multiple TPCDB databases in your storage environment.

After you add multiple connections to the TPCDB data source, when you run a report you select the TPCDB database to connect to.

Procedure

1. Go to the URL for your Cognos Analytics server. The format of the URL is similar to this URL: `http://myhostname:9300/bi`
2. In the Welcome portal, click **Manage > Administration** console.
3. In the Configuration tab, click Data Source Connections.
4. Click the TPCDB data source.
5. On the toolbar, click the New Connection icon .
6. On the Specify a name and description page, type a name for the new connection in the Name field and then click Next.
Do not use TPCDB as the name for the new connection because that conflicts with the existing TPCDB connection.
7. On the Specify the connection page, select IBM DB2 from the Type list.
8. Verify that Use the default object gateway is selected.
9. Clear Configure JDBC connection then click Next.
10. On the Specify the IBM DB2 connection string page, in the DB2 database name property, enter the alias value that you used when you cataloged the remote database.
For example, enter the value `remDB`.
11. In the Signon section, do the following:
 - a. Verify that Signons is selected.
 - b. Click Password.
 - c. In the User ID property, enter the user name of the owner of the DB2® instance where the IBM Spectrum Control database is located.
On Windows operating systems, this user name is typically `db2admin`. On Linux® or AIX® operating systems, this user name is typically `db2inst1`.
 - d. In the Password and Confirm password properties, enter the password that is associated with the User ID property in [11.c](#).
12. In the Testing section, click Test the connection.
13. On the Test the connection page, click Test.
14. When the test is successful, click Close on the View the results page.
15. On the Test the connection page, click Close.
16. On the Specify the IBM DB2 connection string page, click Next.
17. On the Specify the commands page, click Finish.

18. If you scheduled reports, modify the parameters of all scheduled reports to specify the database against which to run the report.

Related information

- 🔗 [Creating Cognos Analytics data source for the IBM Spectrum Control database](#)
- 🔗 [Schedule Management](#)

Customizing the logo and title of reports


You can change the logo that is displayed in all IBM Spectrum® Control reports and the title of individual reports. Before you modify predefined reports, create backup copies of them.

- [Creating a folder in the Cognos Analytics reporting tool](#)
You can create a folder to store backup copies of the predefined reports that you want to modify or for custom reports that you create.
- [Creating a backup copy of a report](#)
Before you modify a predefined report, create a backup copy of the report.
- [Changing the logo for all reports](#)
You can change the logo that is displayed in the output of all reports. You can change the default logo to your company logo or to another image.
- [Modifying the title of a report](#)
You can change the title that is displayed in the output of a report.

Creating a folder in the Cognos Analytics reporting tool

You can create a folder to store backup copies of the predefined reports that you want to modify or for custom reports that you create.

Procedure

1. In the Cognos® Analytics reporting tool, click Team Content in the Welcome portal.
2. Navigate to the location where you want to create the folder. For example, click IBM Spectrum Control Predefined Reports.
3. Click the New folder icon .
4. Type a name for the folder.
For example, type `Backups` as the name of the folder.

Creating a backup copy of a report

Before you modify a predefined report, create a backup copy of the report.

Procedure

1. In the optional Cognos® Analytics reporting tool, create a folder in which to store the backup copy.
2. Navigate to the report of which you want to create a backup copy.
3. Right-click on the report, then click Copy or move.
4. In the Copy or move window, navigate to the folder that you created in step [1](#).
5. Click Copy to.

Related tasks

- [Creating a folder in the Cognos Analytics reporting tool](#)

Changing the logo for all reports

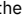
You can change the logo that is displayed in the output of all reports. You can change the default logo to your company logo or to another image.

Before you begin

Make sure that your logo is approximately the same size as the default logo that is in all predefined reports.

Procedure

1. To store images for your company, create a directory, for example, `/myCompany/images`, at the following location:
`Cognos_installation_directory/webcontent/bi/samples`
2. Copy the logo image to the new directory, for example, `../samples/myCompany/images/myLogo.jpg`.
3. In the Welcome portal of the Cognos® Analytics reporting tool, click Team Content > IBM Spectrum Control Report Layouts.

4. Create a backup copy of the report_page_template file.
5. Right-click the report_page_template file, then click Edit report.
6. Click the default image that is displayed underneath the "Double-click to edit text" of the workspace.
7. Click the More icon  for the report, then click Edit Image URL.
8. Modify the image URL to refer to your logo image.
For example, you can specify the following image URL:

`../samples/myCompany/images/myLogo.jpg`

9. Click Save.


Related tasks

- [Creating a backup copy of a report](#)

Modifying the title of a report

You can change the title that is displayed in the output of a report.

Procedure

1. In the Cognos® Analytics reporting tool, navigate to the location of the report whose title you want to change.
2. Right-click the report, then click Properties.
3. Point to the title of the report in the area that displays the properties of the report, then click the Edit icon .
4. Edit the title of the report.

Related tasks

- [Creating a backup copy of a report](#)

Predefined reports about resource relationships

You can run predefined reports about the end-to-end relationship from a server to back-end storage. You can also run reports about the end-to-end relationship from file systems on servers to volumes on storage pools.

- [Running the Storage Resource Relationships Summary report](#)
To show the relationships between storage resources, run the Storage Resource Relationships Summary report.
- [Storage Resource Relationships Summary report](#)
Shows the end-to-end relationship from a server to back-end storage. The report shows servers, hypervisors, NAS, storage systems, storage virtualizers, pools, and volumes.
- [Running the Storage Resource Relationships Summary \(Configurable\) report](#)
To show the relationships between resources, run the Storage Resource Relationships Summary (Configurable) report. You can show or hide resources, and you can hide resources that are not related to other resources or storage devices.
- [Storage Resource Relationships Summary \(Configurable\) report](#)
Shows the end-to-end relationship from a server to back-end storage. You can configure which resources to show in the report.
- [Running the File System to Volume Relationships report](#)
To show the end-to-end relationship from file systems on servers to volumes on storage pools, run the File System to Volume Relationships.
- [File System to Volume Relationships report](#)
Shows the end-to-end relationship from file systems on servers to volumes on storage pools. The report shows the usage of space on file systems over a specified period. Use the report to analyze changes in your file system requirements.

Running the Storage Resource Relationships Summary report

To show the relationships between storage resources, run the Storage Resource Relationships Summary report.

Before you begin

To run the report with the default settings, click Finish.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Resource Relationships Summary.
3. Optional: Select one or more of the following:
 - a. Servers
 - b. Hypervisors
 - c. NAS resources

- d. Storage virtualizers
 - e. Storage systems
4. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.
Restriction: To create a report when you select a single resource, the selected resource must be connected to other resources.

Related reference

- [Storage Resource Relationships Summary report](#)

Storage Resource Relationships Summary report

Shows the end-to-end relationship from a server to back-end storage. The report shows servers, hypervisors, NAS, storage systems, storage virtualizers, pools, and volumes.

Report output

Information is provided about related resources and the devices, such as managed disks (MDisks), that are related to resources.

Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Hypervisor Name

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

NAS Name

The fully qualified domain name of the NAS device. For example, the name of a NAS device might be `nas.example.com`.

Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Volume Name

The name that was assigned to the storage virtualizer volume when it was added to the system.

Storage Virtualizer Volume Assigned Host Connection

The host connection to which the storage virtualizer volume is assigned. The host connection is a definition in the storage virtualizer that contains the WWPN for the server. The storage virtualizer uses the WWPN to assign volumes to servers.

Storage Virtualizer Pool Name

The name that was assigned to the pool when it was added to the system.

Storage Virtualizer MDisk Name

The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum Control shows the name that was defined when the storage system was added for monitoring.

Storage Volume Assigned Host Connection

The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Volume	Performance of One Volume
Managed disk	Performance of One Managed Disk
Storage pool	Most Active Pools, Performance of One Pool

Related tasks

- [Running the Storage Resource Relationships Summary report](#)

Running the Storage Resource Relationships Summary (Configurable) report

To show the relationships between resources, run the Storage Resource Relationships Summary (Configurable) report. You can show or hide resources, and you can hide resources that are not related to other resources or storage devices.

Before you begin

To run the report with the default settings, click Finish. Servers, storage systems, storage virtualizers, and their related resources are shown in the report.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

Tip: To create a report for a single resource, the selected resource must be connected to other storage devices or resources. To create a report about two or more selected resources, the resources that you select must be related.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Resource Relationships Summary (Configurable).
3. Optional: Select the storage resources that you want to show in the report.
4. Optional: Specify whether you want to show or hide a resource that is related to other storage devices such as a file system, but not to other resources.
For example, a server has a related file system. You can show the server and its related storage devices in the report, or you can hide the resource and its related storage devices in the report.
5. Optional: Click Finish to run the report.
6. Optional: Click Next.
In the Select storage resources page, select one or more of the following:
 - a. Servers
 - b. Hypervisors
 - c. NAS resources
 - d. Storage virtualizers
 - e. Storage systems
7. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

Related reference

- [Storage Resource Relationships Summary \(Configurable\) report](#)

Storage Resource Relationships Summary (Configurable) report

Shows the end-to-end relationship from a server to back-end storage. You can configure which resources to show in the report.

Report output

Information is provided about related resources and the devices, such as managed disks (MDisks), that are related to resources.

Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Server Disk Name

The identifier that specifies a path to a server disk, for example `/dev/hdisk0`.

Logical Volume Path

The path to a logical volume on a resource, for example `/dev/hd1`.

File System Mount Point

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be `c:\` or `d:\`. On operating systems such as AIX®, Linux®, Solaris, or HP-UX, the mount point might be `/opt` or `/export/home`.

Hypervisor Name

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

Hypervisor Disk Name

The path that the operating system uses for the server disk.

NAS Name

The fully qualified domain name of the NAS device. For example, the name of a NAS device might be `nas.example.com`.

NAS Export Name

The name of the exported file system.

Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Volume Assigned Host Connection

The host connection to which the storage virtualizer volume is assigned. The host connection is a definition in the storage virtualizer that contains the WWPN for the server. The storage virtualizer uses the WWPN to assign volumes to servers.

Storage Virtualizer Volume Name

The name that was assigned to the storage virtualizer volume when it was added to the system.

Storage Virtualizer Pool Name

The name that was assigned to the pool when it was added to the system.

Storage Virtualizer MDisk Name

The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage Volume Assigned Host Connection

The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Volume	Performance of One Volume
Managed disk	Performance of One Managed Disk
Storage pool	Most Active Pools, Performance of One Pool

Related tasks

- [Running the Storage Resource Relationships Summary \(Configurable\) report](#)

Running the File System to Volume Relationships report

To show the end-to-end relationship from file systems on servers to volumes on storage pools, run the File System to Volume Relationships.

Before you begin

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click File System to Volume Relationships.
3. Select one or more servers.
4. Select the storage virtualizers and storage systems that the storage pools are in.
You can select both storage virtualizers and storage systems, storage virtualizers only, or storage systems only.
5. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

Restriction: To create a report when you select a single resource, the selected resource must be connected to other resources.

Related reference

- [File System to Volume Relationships report](#)

File System to Volume Relationships report

Shows the end-to-end relationship from file systems on servers to volumes on storage pools. The report shows the usage of space on file systems over a specified period. Use the report to analyze changes in your file system requirements.

Report output

The information that is provided about the servers, storage virtualizers, storage systems, pools, and volumes is as follows:

Server Name

The fully qualified domain name of the server. For example, the name of a server might be **server.example.com**.

File System Mount Point

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be c:\ or d:\. On operating systems such as AIX®, Linux®, Solaris, or HP-UX, the mount point might be /opt or /export/home.

Server Disk Name

The identifier that specifies a path to a server disk, for example `/dev/hdisk0`.

File System Capacity (GiB)
The amount of storage space on the file system of the resource.

File System Available Space (GiB)
The amount of unused storage space in the file system of the resource.

File System Available Space Percentage
The percentage of space in the file system that is unused.

Storage Virtualizer Name
A user-defined name of the storage virtualizer. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage virtualizer was added for monitoring.

Storage Virtualizer Volume Name
The name that was assigned to the storage virtualizer volume when it was added to the system.

Storage Virtualizer Volume Is Thin Provisioned
Shows whether a volume or volume copy is thin-provisioned. If this value is **yes**, the resource is thin-provisioned.

Storage Virtualizer Volume Allocated Space Percentage
The percentage of space on a volume that is reserved.
This property applies only to volumes on resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Volume Used Space Percentage
The percentage of reserved space that is being used on a volume.
This property applies only to volumes on resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Pool Name
The name that was assigned to the pool when the pool was added to the system.

Storage Virtualizer Pool Available Space (GiB)
The amount of space in a pool that is not reserved for volumes.

Storage System Name
A user-defined name of the storage system. If a name was not defined, IBM Spectrum Control shows the name that was defined when the storage system was added for monitoring.

Storage Volume Name
The name that was assigned to the storage volume when it was added to the system.

Storage Volume Is Thin Provisioned
Shows whether a volume or volume copy is thin-provisioned. If this value is **yes**, the resource is thin-provisioned.

Storage Volume Allocated Space Percentage
The percentage of space on a volume that is reserved.
This property applies only to volumes on SAN Volume Controller and Storwize® V7000 systems.

Storage Volume Used Space Percentage
The percentage of reserved space that is being used on a volume.
This property applies only to volumes on SAN Volume Controller and Storwize V7000 systems.

Storage Pool Name
The name that was assigned to the pool when it was added to the system.

Storage Pool Available Space (GiB)
The amount of unused space that is not reserved for volumes in pools that are on the storage system.
IBM Spectrum Control uses the following formula to determine this value:

pool capacity - used space

For XIV® systems, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Related tasks

- [Running the File System to Volume Relationships report](#)

Predefined reports about switches and switch ports

You can run predefined reports to analyze and compare the performance of switches and switch ports.

- [Predefined reports about switches](#)
You can run predefined reports to analyze and compare the performance of switches, and to see information about the ports that are available from switches.
- [Predefined reports about switch ports](#)
You can run predefined reports to analyze and compare the performance of switch ports.

Predefined reports about switches

You can run predefined reports to analyze and compare the performance of switches, and to see information about the ports that are available from switches.

- [Running the Availability of Switch Ports report](#)
To see information about the ports and connected ports that are available from switches, run the Availability of Switch Ports report.
- [Availability of Switch Ports report](#)
Shows a chart of the 20 switches with the most ports and connected ports, and a table with the properties of the switches.
- [Running the Performance of One Switch report](#)
To see information about the performance of a switch, run the Performance of One Switch report.
- [Performance of One Switch report](#)
Shows three charts, and a table of performance metrics for a switch over a time period that you specify. Use the report to analyze the performance of a switch.
- [Running the Compare Performance of Multiple Switches report](#)
Use the Compare Performance of Multiple Switches report to compare up to four performance metrics for multiple switches. For example, you can compare

bandwidth percentages, data rates, and frame rates.

- [Compare Performance of Multiple Switches report](#)

Shows up to four performance metrics for multiple switches over time. Use the report to compare the performance metrics for switches.

- [Running the Compare Performance of One Switch over Time Ranges report](#)

To compare a performance metric for one switch in two date ranges, run the Compare Performance of One Switch over Time Ranges report.

- [Compare Performance of One Switch over Time Ranges report](#)

Shows one performance metric on one switch over two time periods. For example, you can use the report to compare the I/O rate for this week and last week.

Running the Availability of Switch Ports report

To see information about the ports and connected ports that are available from switches, run the Availability of Switch Ports report.

Before you begin

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

Procedure

1. In the Welcome portal, click Team Content, IBM Spectrum Control Predefined Reports.
2. Click Fabrics and Switches.
3. Click Switches.
4. Click Availability of Switch Ports.
5. Select the switches, and then select a sort order.
6. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Availability of Switch Ports report](#)

Availability of Switch Ports report

Shows a chart of the 20 switches with the most ports and connected ports, and a table with the properties of the switches.

Charts

The chart shows the number of ports for each switch that are not online, and the number of ports that are online and connected to other ports.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

For each switch, the following information is provided:

Switch Name

The logical name of the switch where a Fibre Channel port is located, or the name that was defined when the switch was added for monitoring. If neither name is available, IBM Spectrum® Control uses the WWN of the switch.

Switch Location

The physical location of the switch. The location is defined when a switch is added to IBM Spectrum Control. You can add or edit the location of the switch on the General tab of the properties notebook.

Switch WWN

The World Wide Name (WWN) of the switch. A WWN is the unique 64-bit identifier for the switch.

Switch Is Virtual

Shows whether a switch is a logical switch. If this value is **Yes**, the switch is a logical switch.

Switch IP Address

The IP address of the resource.

Switch Vendor

The vendor who supplied the resource.

Switch Model

The model name or model number of the resource.

Switch Ports

The total number of ports on the switch that are not online, and ports that are online and connected to other ports.

Switch Connected Ports

The number of ports that are connected to a storage resource, where the storage resource can be a storage system, server, or hypervisor. This value is only available when a switch is viewed as a related resource of a storage system, server, or hypervisor.

Switch Last Data Collection

The date and time when storage statistics were last collected from the resource.
Switch Last Data Collection Status
The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

Related tasks

- [Running the Availability of Switch Ports report](#)

Running the Performance of One Switch report

To see information about the performance of a switch, run the Performance of One Switch report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Fabrics and Switches.
3. Click Switches.
4. Click Performance of One Switch.
5. Select a switch.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One Switch report](#)

Performance of One Switch report

Shows three charts, and a table of performance metrics for a switch over a time period that you specify. Use the report to analyze the performance of a switch.

Charts

The charts show the following performance metrics for the ports on the switch:

- The send frame rates, the receive frame rates, and the total frame rates
- The send data rates, the receive data rates, and the total data rates
- The send bandwidth percentages, the receive bandwidth percentages, and the overall bandwidth percentages

Report output

The following information is provided in the report for the period that you specified:

Switch Name

The logical name of the switch or the name that was defined when the switch was added to IBM Spectrum® Control. If neither name is available, the WWN of the switch is displayed.

Switch Location

The physical location of the switch. The location is defined when a switch is added to IBM Spectrum Control. You can add or edit the location of the switch on the General tab of the properties notebook.

Port Send Frame Rate (frames/s)

The average number of frames per second that are sent by the port.

Port Receive Frame Rate (frames/s)

The average number of frames per second that are received by the port.

Total Port Frame Rate (frames/s)

The average number of frames per second that are transferred. This value includes frames that are sent and received by the port.

Port Send Data Rate (MiB/s)

The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)

The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Bandwidth Percentage

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage

The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Related reports

To open the Most Active Switch Ports report, click the name of the switch in the report table.

Related tasks

- [Running the Performance of One Switch report](#)

Running the Compare Performance of Multiple Switches report

Use the Compare Performance of Multiple Switches report to compare up to four performance metrics for multiple switches. For example, you can compare bandwidth percentages, data rates, and frame rates.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Fabrics and Switches.
3. Click Switches.
4. Click Compare Performance of Multiple Switches.
5. Select the switches, and then click Next.
6. Select the category of performance metric that you want to include in the report, and then select the performance metric.
7. Optional: To add more metrics to the report, complete the following steps:
 - a. Click Add Metric.
 - b. Select the category of performance metric, and then select the performance metric.
 - c. If you want to remove the last metric that you added, click Remove Metric.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period.

It is easier to analyze small amounts of information over shorter periods.

Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
10. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple Switches report](#)

Compare Performance of Multiple Switches report

Shows up to four performance metrics for multiple switches over time. Use the report to compare the performance metrics for switches.

Charts

The charts show the performance metrics for the switches that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metrics

Shows the data for each performance metric.

Switch names

Shows the names of the switches, and the performance metric data for each switch.

Related tasks

- [Running the Compare Performance of Multiple Switches report](#)

Running the Compare Performance of One Switch over Time Ranges report

To compare a performance metric for one switch in two date ranges, run the Compare Performance of One Switch over Time Ranges report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Fabrics and Switches.
3. Click Switches.
4. Click Compare Performance of One Switch over Time Ranges.
5. Select the switch.
6. Select the category of performance metric about which you want a report, and then select the performance metric.
7. Select an interval.
8. Specify the two date ranges or reporting periods that you want to compare.
For example, you can compare the performance metric for this week with last week.
Tip: To specify a start date and an end date for a reporting period, select Custom Date Range from the list of reporting periods.
9. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Compare Performance of One Switch over Time Ranges report](#)

Compare Performance of One Switch over Time Ranges report

Shows one performance metric on one switch over two time periods. For example, you can use the report to compare the I/O rate for this week and last week.

Charts

The chart shows the performance data of the switch on the performance metric you selected, over the time periods that you selected.

Report output

The following information is provided in the report:

Switch Name

The logical name of the switch or the name that was defined when the switch was added to IBM Spectrum® Control. If neither name is available, the WWN of the switch is displayed.

Day in Range

The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The number of days that are shown is determined by the reporting period that was selected for the report. If Last 30 Days was selected, then entries for days 1 - 30 are shown on the report.

Hour in Range

The number of the hour in the day at which the performance metric data was collected. For example, 0 is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

Minute in Range

The interval of minutes during which the performance metric data was collected. For example, 0 shows that performance data was collected 0 - 5 minutes after the hour. Similarly, 5 shows that the performance data that was collected 5 - 10 minutes after the hour.

First Date Range

The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the first date range.

Second Date Range

The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the second date range.

Related tasks

- [Running the Compare Performance of One Switch over Time Ranges report](#)

Predefined reports about switch ports

You can run predefined reports to analyze and compare the performance of switch ports.

- [Running the Most Active Switch Ports report](#)
To see which ports are most active on switches, run the Most Active Switch Ports report.
- [Most Active Switch Ports report](#)
Shows a chart of the 20 switches that are most active, and details for all switches, for a time period that you specify. Use the report to analyze the performance of switches.
- [Running the Performance of One Switch Port report](#)
To see information about the performance of a switch port, run the Performance of One Switch Port report.
- [Performance of One Switch Port report](#)
Shows four charts, and a table of performance metrics for a switch port. Use this report to view multiple performance metric types on the same chart. For example, you can view I/O rate and response time metrics on the same chart.
- [Running the Compare Performance of Multiple Switch Ports report](#)
Use the Compare Performance of Multiple Switch Ports report to compare up to four performance metrics for multiple switch ports. For example, you can compare bandwidth percentages, data rates, and I/O rates.
- [Compare Performance of Multiple Switch Ports report](#)
Shows up to four performance metrics for multiple switch ports over time. Use the report to compare the performance metrics for switch ports.
- [Running the Compare Performance of One Switch Port over Time Ranges report](#)
To compare a performance metric for one switch port in two date ranges, run the Compare Performance of One Switch Port over Time Ranges report.
- [Compare Performance of One Switch Port over Time Ranges report](#)
Shows one performance metric on one switch port over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Running the Most Active Switch Ports report

To see which ports are most active on switches, run the Most Active Switch Ports report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Fabrics and Switches.
3. Click Switch Ports.
4. Click Most Active Switch Ports.
5. Select the switches, and then select a sort order.
6. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
7. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Switch Ports report](#)

Most Active Switch Ports report

Shows a chart of the 20 switches that are most active, and details for all switches, for a time period that you specify. Use the report to analyze the performance of switches.

Charts

The bar chart shows the most active ports on the switch. By default, ports are sorted by the total frame rate.
Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active ports on switches is as follows:

Switch Name

The logical name of the switch or the name that was defined when the switch was added to IBM Spectrum® Control. If neither name is available, the WWN of the switch is displayed.

Switch Blade Slot Number

The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

Switch Port Number

The number of the port on the switch.

Switch Port WWPN

The worldwide port name of the port on the switch.

Total Port Frame Rate (frames/s)
The average number of frames per second that are transferred. This value includes frames that are sent and received by the port.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Switch	Performance of One Switch
Switch port	Performance of One Switch Port

Related tasks

- [Running the Most Active Switch Ports report](#)

Running the Performance of One Switch Port report

To see information about the performance of a switch port, run the Performance of One Switch Port report.

Before you begin

You must select a switch and a switch port. If the port is on a blade, you must select the number of the slot on the switch to which the blade is attached.

Procedure

1. In the Welcome portal, click Team Content, IBM Spectrum Control Predefined Reports.
2. Click Fabrics and Switches.
3. Click Switch Ports.
4. Click Performance of One Switch Port.
5. Select a switch.
6. Optional: Select the number of the blade slot.
If the port is on a blade, you must select the number of the blade slot.
7. Select a port.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
10. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One Switch Port report](#)

Performance of One Switch Port report

Shows four charts, and a table of performance metrics for a switch port. Use this report to view multiple performance metric types on the same chart. For example, you can view I/O rate and response time metrics on the same chart.

Charts

The charts show the following performance statistics for the port:

- The send frame rate, the receive frame rate, and the total frame rate
- The send data rate, the receive data rate, and the total data rate
- The send bandwidth percentage, the receive bandwidth percentage, and the overall bandwidth percentage
- The error rates

Report output

The following information is provided in the report for the period that you specified:

Switch Name

The logical name of the switch or the name that was defined when the switch was added to IBM Spectrum® Control. If neither name is available, the WWN of the switch is displayed.

Switch Blade Slot Number

The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

Switch Port Number

The number of the port on the switch.

Switch Port WWPN

The worldwide port name of the port on the switch.

Port Send Frame Rate (frames/s)

The average number of frames per second that are sent by the port.

Port Receive Frame Rate (frames/s)

The average number of frames per second that are received by the port.

Total Port Frame Rate (frames/s)

The average number of frames per second that are transferred. This value includes frames that are sent and received by the port.

Port Send Data Rate (MiB/s)

The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)

The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Bandwidth Percentage

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage

The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Error Frame Rate (count/s)

The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

Discarded Frame Rate (count/s)

The average number of frames per second that are discarded because host buffers are unavailable for the port.

Link Failure Rate (count/s)

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

CRC Error Rate (count/s)

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.

Discarded Class 3 Frame Rate (count/s)

The average number of class 3 frames per second that are discarded.

Invalid Transmission Word Rate (count/s)

The average number of bit errors per second that are detected.

Related reports

To open the Most Active Switch Ports report, click the name of the switch in the report table.

Related tasks

- [Running the Performance of One Switch Port report](#)

Running the Compare Performance of Multiple Switch Ports report

Use the Compare Performance of Multiple Switch Ports report to compare up to four performance metrics for multiple switch ports. For example, you can compare bandwidth percentages, data rates, and I/O rates.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Fabrics and Switches.
3. Click Switch Ports.
4. Click Compare Performance of Multiple Switch Ports.
5. Select the switches.
6. Select switch ports. To select switch ports, complete the following steps:
 - a. Type a keyword.
For example, type `example1` to find the names of switch ports that start with `example1`.
To change the default search option, click Options.
Use the percent sign as a wildcard.

- For example, type a% to find switch ports with names begin with 'a' or 'A'. Type %a to find switch ports with names that begin with or contain the letter 'a' or 'A'. You can type % to retrieve all of the switch ports on a resource.
- b. Click Search.
If the search is successful, the switch ports are displayed in the Results list.
 - c. Select the switch ports that you require from the list, and then click Insert.
7. Click Next.
 8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
 9. Optional: To add more metrics to the report, complete the following steps:
 - a. Click Add Metric.
 - b. Select the category of performance metric, and then select the performance metric.
 - c. If you want to remove the last metric that you added, click Remove Metric.
 10. Optional: Select an interval.
 11. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
 12. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple Switch Ports report](#)

Compare Performance of Multiple Switch Ports report

Shows up to four performance metrics for multiple switch ports over time. Use the report to compare the performance metrics for switch ports.

Charts

The charts show the performance metrics for the switch ports that you selected for the period that you specified.

Report output

The following information is provided in the report:

- Performance metrics
Shows the data for each performance metric.
- Switch names
Shows the names of the switches, and the performance metric data for each switch.
- Slot numbers
Shows the slot numbers, and the performance metric data for each slot.
- Switch port numbers
Shows the port numbers, and the performance metric data for each switch port number.

Related tasks

- [Running the Compare Performance of Multiple Switch Ports report](#)

Running the Compare Performance of One Switch Port over Time Ranges report

To compare a performance metric for one switch port in two date ranges, run the Compare Performance of One Switch Port over Time Ranges report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Fabrics and Switches.
3. Click Switch Ports.
4. Click Compare Performance of One Switch Port over Time Ranges.
5. Select the switch that the port is on.
6. If the switch is on a blade, select the slot number of the blade that the port is on.
7. Select the number of the port.
8. Select the category of performance metric about which you want a report, and then select the performance metric.
9. Select an interval.
10. Specify the two date ranges or reporting periods that you want to compare.
For example, you can compare the performance metric for this week with last week.
Tip: To specify a start date and an end date for a reporting period, select Custom Date Range from the list of reporting periods.
11. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Compare Performance of One Switch Port over Time Ranges report](#)

Compare Performance of One Switch Port over Time Ranges report

Shows one performance metric on one switch port over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Charts

The chart shows the performance data of the switch port on the performance metric you selected, over the time periods that you selected.

Report output

The following information is provided in the report:

Switch Name

The logical name of the switch or the name that was defined when the switch was added to IBM Spectrum® Control. If neither name is available, the WWN of the switch is displayed.

Switch Blade Slot Number

The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

Switch Port Number

The number of the port on the switch.

Switch Port WWPN

The worldwide port name of the port on the switch.

Day in Range

The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The number of days that are shown is determined by the reporting period that was selected for the report. If Last 30 Days was selected, then entries for days 1 - 30 are shown on the report.

Hour in Range

The number of the hour in the day at which the performance metric data was collected. For example, 0 is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

Minute in Range

The interval of minutes during which the performance metric data was collected. For example, 0 shows that performance data was collected 0 - 5 minutes after the hour. Similarly, 5 shows that the performance data that was collected 5 - 10 minutes after the hour.

First Date Range

The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the first date range.

Second Date Range

The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the second date range.

Related tasks

- [Running the Compare Performance of One Switch Port over Time Ranges report](#)

Predefined reports about groups

You can run predefined reports to analyze the capacity of groups. A group is a set of logically related volumes, file systems, and shares. For example, a group that represents a business critical application might include the volumes, file systems, and shares that provide storage to the application.

- [Running the Capacity of One Group report](#)
To see information about the block-level and file-level capacity and space of a group, run the Capacity of One Group report.
- [Capacity of One Group report](#)
Shows charts and tables of the capacity, available space, and used space statistics for each type of resource in a group. The report shows separate charts and tables for block-level and file-level resources.
- [Running the Groups Capacity report](#)
To see information about the capacity and space of groups, run the Groups Capacity report.
- [Groups Capacity report](#)
Shows a chart of the 20 groups with the most capacity, available space, or used space, and a table with details of all groups. Use the report to review the space allocation in your groups.

Running the Capacity of One Group report

To see information about the block-level and file-level capacity and space of a group, run the Capacity of One Group report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Groups.
3. Click Capacity of One Group.
4. Select a group.
5. Optional: Select a sort order.
6. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Capacity of One Group report](#)

Capacity of One Group report

Shows charts and tables of the capacity, available space, and used space statistics for each type of resource in a group. The report shows separate charts and tables for block-level and file-level resources.

Charts

The charts show the following block-level and file-level capacity statistics for a group:

- Total capacity
- Total available space
- Total used space

If there are any block-level storage devices in the group, the chart also shows the total allocated space for those devices.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

The report provides the following information about the group:

Group Name

The name that was assigned to the group by the user when the group was created.

Resource Name

For a server or NAS file server, this property shows the fully qualified domain name of the server or NAS file server.

For a block-level storage resource, this property shows the user-defined name of the resource.

Resource Type

For a server or NAS file server, this property shows the operating system that is running on the server or NAS file server.

For a block-level storage resource, this property shows the type of resource. For example, the property can be IBM System Storage DS8800 or another type of resource.

Total File Capacity (GiB)

The total amount of file-level storage space that is on the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

Total File Used Space (GiB)

The total amount of file-level storage space that is used by the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

Total File Available Space (GiB)

The total amount of file-level unused space that is on the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

Total Block Capacity (GiB)

The total amount of block-level storage space that is committed to storage system volumes or storage virtualizer volumes in a group.

Total Block Used Space (GiB)

The total amount of block-level allocated space that is used by storage system volumes or storage virtualizer volumes in a group.

Total Block Available Space (GiB)

The total amount of block-level unused space that is available on storage system volumes or storage virtualizer volumes in a group.

Total Block Allocated Space (GiB)

The total amount of block-level space that is reserved for the storage system volumes or storage virtualizer volumes in a group. The allocated space includes both thin-provisioned and standard volumes.

Related tasks

- [Running the Capacity of One Group report](#)

Running the Groups Capacity report

To see information about the capacity and space of groups, run the Groups Capacity report.

Before you begin

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Groups.
3. Click Groups Capacity.
4. Optional: Select one or more groups.
5. Optional: Select a sort order.
6. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Groups Capacity report](#)

Groups Capacity report

Shows a chart of the 20 groups with the most capacity, available space, or used space, and a table with details of all groups. Use the report to review the space allocation in your groups.

Charts

The bar chart shows the following block-level and file-level capacity statistics for groups:

- Total available space
- Total used space

Tip: To change the sort order, select a value from the Sort Order list.

Report output

For each group, the following information is provided:

Group Custom Tag 1

User-defined text that is associated with a group.

Group Name

The name that was assigned to the group by the user when the group was created.

Total Capacity (GiB)

The total amount of block-level and file-level storage space that is committed to volumes, file systems, and shares.

Total Used Space (GiB)

The total amount of block-level and file-level space that is used by volumes, file systems, and shares.

Total Available Space (GiB)

The total amount of block-level and file-level space that is available on volumes, file systems, and shares.

Total Block Capacity (GiB)

The total amount of block-level storage space that is committed to storage system volumes or storage virtualizer volumes in a group.

Total Block Used Space (GiB)

The total amount of block-level allocated space that is used by storage system volumes or storage virtualizer volumes in a group.

Total Block Available Space (GiB)

The total amount of block-level unused space that is available on storage system volumes or storage virtualizer volumes in a group.

Total Block Allocated Space (GiB)

The total amount of block-level space that is reserved for the storage system volumes or storage virtualizer volumes in a group. The allocated space includes both thin-provisioned and standard volumes.

Total File Capacity (GiB)

The total amount of file-level storage space that is on the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

Total File Used Space (GiB)

The total amount of file-level storage space that is used by the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

Total File Available Space (GiB)

The total amount of file-level unused space that is on the file systems in the group. The space is exported space, that is, the space is on file systems that are made available to remote clients over a network.

Related tasks

- [Running the Groups Capacity report](#)

Predefined reports about hypervisors

You can run predefined reports to analyze the capacity and performance of hypervisors.

- [Running the Hypervisors Capacity report](#)
To see information about the capacity of disks on hypervisors, run the Hypervisors Capacity report.
- [Hypervisors Capacity report](#)
Shows a chart of the 20 hypervisors with the most allocated space, and a table with space statistics for all hypervisors. Use the report to review space allocation in your hypervisors.
- [Running the Most Active Hypervisors report](#)
To see which hypervisors are the most active, run the Most Active Hypervisors report.
- [Most Active Hypervisors report](#)
Shows a chart of the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load.
- [Running the Summarized Performance of Volumes by Hypervisor report](#)
To see the load that a particular hypervisor adds to your storage environment, use the Summarized Performance of Volumes by Hypervisor report.
- [Summarized Performance of Volumes by Hypervisor report](#)
Shows summarized performance metrics for volumes on a hypervisor. Use the report to identify the aggregate load and the average response time of a particular hypervisor.
- [Running the Performance of Volumes by Hypervisor report](#)
To see performance metrics for the volumes on a hypervisor, use the Performance of Volumes by Hypervisor report.
- [Performance of Volumes by Hypervisor report](#)
Shows the performance metrics for volumes on a hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.
- [Running the Hypervisor Data Stores Capacity report](#)
To see information about the capacity of data stores on hypervisors, run the Hypervisor Data Stores Capacity report.
- [Hypervisor Data Stores Capacity report](#)
Shows a chart of the 20 hypervisor data stores with the most allocated space. This report also shows a table with space statistics for all hypervisor data stores. Use the report to review space allocation in your hypervisor data stores.
- [Running the Hypervisor Disks Capacity report](#)
To see information about the disk capacity on hypervisors, run the Hypervisor Disks Capacity report.
- [Hypervisor Disks Capacity report](#)
Shows a chart of the 20 hypervisor disks with the most allocated space, and a table with space statistics for all hypervisor disks. Use the report to review space allocation in your hypervisor disks.

Running the Hypervisors Capacity report

To see information about the capacity of disks on hypervisors, run the Hypervisors Capacity report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Hypervisors.
3. Click Hypervisors Capacity.
4. Optional: Select a hypervisor cluster.
5. Select one or more hypervisors.
6. Optional: Select a storage resources sort order.
7. Optional: Select a sort order.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Hypervisors Capacity report](#)

Hypervisors Capacity report

Shows a chart of the 20 hypervisors with the most allocated space, and a table with space statistics for all hypervisors. Use the report to review space allocation in your hypervisors.

Charts

The bar chart shows the used space and available space on hypervisors.

Report output

For each hypervisor, the following information is provided:

Hypervisor Name

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

Hypervisor Location

The physical location of a hypervisor. The location is defined when a hypervisor is added to IBM Spectrum® Control. You can add or edit the location of the hypervisor in the Properties pane of the hypervisor.

Hypervisor Cluster Name

The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

Hypervisor Total Disk Capacity (GiB)

The amount of storage space that is on a hypervisor.

Hypervisor Total Disk Available Space (GiB)

The amount of unused storage space on all disks on a hypervisor.

Hypervisor Total Disk Used Space (GiB)

The amount of storage space that is used on all disks on a hypervisor.

Hypervisor Total Disk Used Space (%)

The amount of storage space that is used on a hypervisor.

Hypervisor Data Store Capacity (GiB)

The amount of storage space that is on the hypervisor data store.

Hypervisor Data Store Available Space (GiB)

The amount of unused storage space on a hypervisor data store.

Hypervisor Data Store Used Space (GiB)

The amount of storage space that is used on a hypervisor data store.

Hypervisor Data Store Used Space (%)

The percentage of storage space that is used on all disks on a hypervisor.

Related reports

To open related reports, click the name of the hypervisor and then click the name of the report. You can open the following related reports:

- Hypervisor Disks Capacity
- Hypervisor Data Stores Capacity

Related tasks

- [Running the Hypervisors Capacity report](#)

Running the Most Active Hypervisors report

To see which hypervisors are the most active, run the Most Active Hypervisors report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Hypervisors.
3. Click Most Active Hypervisors.
4. Select the hypervisors, and then select a sort order.
5. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
6. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Hypervisors report](#)

Most Active Hypervisors report

Shows a chart of the 20 hypervisors that have the most active volumes. The most active volumes have the greatest aggregate load.

Charts

The bar chart shows the most active hypervisors. By default, hypervisors are sorted by the total I/O rate.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active hypervisors is as follows:

Hypervisor Name

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be **hypervisor.example.com**.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Volume Utilization

The average percentage of time that the volume is busy.

Related reports

To open the Summarized Performance of Volumes by Hypervisor report, click the name of the hypervisor in the report table.

Related tasks

- [Running the Most Active Hypervisors report](#)

Running the Summarized Performance of Volumes by Hypervisor report

To see the load that a particular hypervisor adds to your storage environment, use the Summarized Performance of Volumes by Hypervisor report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Hypervisors.
3. Click Summarized Performance of Volumes by Hypervisor.
4. Select one or more hypervisors.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Summarized Performance of Volumes by Hypervisor report](#)

Summarized Performance of Volumes by Hypervisor report

Shows summarized performance metrics for volumes on a hypervisor. Use the report to identify the aggregate load and the average response time of a particular hypervisor.

Charts

The charts show the following performance metrics for the volumes on a hypervisor:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, the write data rates, and the total data rates
- The read cache percentages and the write cache percentages

Report output

The following information is provided in the report for the period that you specified:

Hypervisor Name

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be **hypervisor.example.com**.

Day, hour, or time

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation, across all volumes on the hypervisor. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second, across all volumes on the hypervisor. This value includes both sequential and nonsequential operations. The average value for each volume is added to give a total I/O rate for all volumes on the hypervisor.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations, across all volumes on the hypervisor. The average value for each volume is added to give a total data rate for all volumes on the hypervisor.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation, across all volumes on the hypervisor.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation, across all volumes on the hypervisor.

Read Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations, across all volumes on the hypervisor. The average value for each volume is added to give a read data rate for all volumes on the hypervisor.

Write Data Rate (MiB/s)

The average number of MiB per second that are transferred for write operations, across all volumes on the hypervisor. The average value for each volume is added to give a write data rate for all volumes on the hypervisor.

Overall Read I/O Rate (ops/s)

The average number of read operations per second, which is totaled for all volumes on the hypervisor. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second, which is totaled for all volumes on the hypervisor. This value includes both sequential and nonsequential write operations.

Overall Read Cache Hit Percentage

The average percentage of all read operations that find data in the cache, across all volumes on the hypervisor. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage

The average percentage of all write operations that are handled in the cache, across all volumes on the hypervisor. This value includes both sequential and nonsequential write operations.

Related reports

To open the Performance of Volumes by Hypervisor report, click the name of the hypervisor in the report table.

Related tasks

- [Running the Summarized Performance of Volumes by Hypervisor report](#)

Running the Performance of Volumes by Hypervisor report

To see performance metrics for the volumes on a hypervisor, use the Performance of Volumes by Hypervisor report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Hypervisors.

3. Click Performance of Volumes by Hypervisor.
4. Select one or more hypervisors.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of Volumes by Hypervisor report](#)

Performance of Volumes by Hypervisor report

Shows the performance metrics for volumes on a hypervisor. Use the report to identify the contribution of individual volumes to the load of a particular hypervisor in your storage environment.

Charts

The charts show the following performance metrics for the volumes:

- The total overall I/O rates. The chart shows the contribution of each volume to the overall load on the hypervisor.
- The total data rates. The chart shows the contribution of each volume to the overall load on the hypervisor.
- The volume utilization percentages. The utilization percentage of each volume is displayed relative to the utilization percentage of the other volumes.
- The total response times. The response time of each volume is displayed relative to the response time of the other volumes.

Report output

The following information is provided in the report for the period that you specified:

Hypervisor Name

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

Day, hour, or time

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Volume Utilization

The average percentage of time that the volume is busy.

Related tasks

- [Running the Performance of Volumes by Hypervisor report](#)

Running the Hypervisor Data Stores Capacity report

To see information about the capacity of data stores on hypervisors, run the Hypervisor Data Stores Capacity report.

Before you begin

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Hypervisors.
3. Click Data Stores Capacity.
4. Optional: Select a hypervisor cluster.
5. Select one or more hypervisors.
6. Optional: Select a sort order.
7. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Hypervisor Data Stores Capacity report](#)

Hypervisor Data Stores Capacity report

Shows a chart of the 20 hypervisor data stores with the most allocated space. This report also shows a table with space statistics for all hypervisor data stores. Use the report to review space allocation in your hypervisor data stores.

Charts

The bar chart shows the used space and available space for data stores on hypervisors.

Report output

For each hypervisor, the following information is provided:

Hypervisor Cluster Name

The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

Hypervisor Data Store

The path to the VMware ESX data store on the hypervisor, for example datastore_svc3c.

Hypervisor Data Store Cluster Name

The name of the cluster, if the data store is a member of a data store cluster. If the data store is not a member of a data store cluster, no information is displayed.

Hypervisor Data Store Type

The type of file system that the hypervisor data store uses. For example, the type can be NFS or VMFS.

Hypervisor System Mount Point

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be c:\ or d:\. On operating systems such as AIX®, Linux®, Solaris, or HP-UX, the mount point might be /opt or /export/home.

Hypervisor Data Store Capacity (GiB)

The amount of storage space that is on the hypervisor data store.

Hypervisor Data Store Used Space (GiB)

The amount of storage space that is used on a hypervisor data store.

Hypervisor Data Store Available Space (GiB)

The amount of unused storage space on a hypervisor data store.

Hypervisor Data Store Used Space (%)

The percentage of storage space that is used on all disks on a hypervisor.

Hypervisor File System Capacity (GiB)

The amount of storage space on the file system of the hypervisor.

Hypervisor File System Available Space (GiB)

The amount of unused storage space in the file system of the hypervisor.

Related tasks

- [Running the Hypervisor Data Stores Capacity report](#)

Running the Hypervisor Disks Capacity report

To see information about the disk capacity on hypervisors, run the Hypervisor Disks Capacity report.

Before you begin

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Hypervisors.
3. Click Hypervisor Disks Capacity.
4. Optional: Select a hypervisor cluster.
5. Select one or more hypervisors.
6. Optional: Select a sort order.
7. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Hypervisor Disks Capacity report](#)

Hypervisor Disks Capacity report

Shows a chart of the 20 hypervisor disks with the most allocated space, and a table with space statistics for all hypervisor disks. Use the report to review space allocation in your hypervisor disks.

Charts

The bar chart shows the used space and available space on hypervisors.

Report output

For each hypervisor, the following information is provided:

Hypervisor Cluster Name

The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

Hypervisor Disk Name

The path that the operating system uses for the server disk.

Hypervisor Disk Capacity (GiB)

The amount of storage space that is on a hypervisor disk.

Hypervisor Disk Available Space (GiB)

The amount of unused storage space that is on a hypervisor disk.

Hypervisor Disk Used Space (GiB)

The amount of storage space that is used on a hypervisor disk.

Hypervisor Disk Used Space (%)

The amount of storage space that is used on a hypervisor disk.

Hypervisor Disk Status

The status of a disk on the hypervisor. Use the status to determine the condition of the disk, and if any actions must be taken. For example, if a disk has an Error status, take immediate action to correct the problem. If the disk has an Operational status, then it is operating normally and no further action is required.

Hypervisor Disk Vendor

The vendor who supplied the hypervisor disk.

Hypervisor Disk Model

The model name or model number of the hypervisor disk.

Hypervisor Disk Serial Number

The serial number of the hypervisor disk.

Hypervisor Disk Firmware Version

The version number of the firmware that is running on the hypervisor disk.

Related tasks

- [Running the Hypervisor Disks Capacity report](#)

Predefined reports about servers and file systems on servers

You can run predefined reports to analyze the capacity and performance of servers, file systems, and volumes.

- [Predefined reports about servers](#)
You can run predefined reports to analyze the capacity and performance of servers and volumes.
- [Predefined reports about file systems on servers](#)
You can run predefined reports to analyze the capacity of file systems that are grouped by server.

Predefined reports about servers

You can run predefined reports to analyze the capacity and performance of servers and volumes.

- [Running the Servers Capacity report](#)
To see information about the capacity and space of server disks and file systems, run the Servers Capacity report.
- [Servers Capacity report](#)
Shows a summary of the available and used disk space, and the available and used file systems, for servers. Use the report to identify servers that are running out of space.
- [Running the Most Active Servers report](#)
To see which servers are the most active, run the Most Active Servers report.
- [Most Active Servers report](#)
Shows a chart of the 20 servers that have the most active volumes. The most active volumes have the greatest aggregate load.
- [Running the Summarized Performance of Volumes by Server report](#)
To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Server report.
- [Summarized Performance of Volumes by Server report](#)
Shows summarized performance metrics for volumes on a server. Use the report to identify the aggregate load and the average response time of a particular server.
- [Running the Performance of Volumes by Server report](#)
To see performance metrics for the volumes on a server, use the Performance of Volumes by Server report.
- [Performance of Volumes by Server report](#)
Shows the performance metrics for volumes on a server. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.
- [Running the Server Disks Capacity report](#)
To see information about the capacity of disks on servers, run the Server Disks Capacity report.
- [Server Disks Capacity report](#)
Shows the available and used disk space for all disks, which are grouped by server. You can use the report to identify disks that have no allocated space.

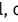
Running the Servers Capacity report

To see information about the capacity and space of server disks and file systems, run the Servers Capacity report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content,  IBM Spectrum Control Predefined Reports.
2. Click Servers.
3. Click Servers Capacity.
4. Optional: Select one or more servers.
5. Optional: Select a sort option.
You can order the metrics represented in the chart by disk or by file system.
Tip: When the report is created, you can change the sort option.
6. Optional: Select a sort order.
7. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Servers Capacity report](#)

Servers Capacity report

Shows a summary of the available and used disk space, and the available and used file systems, for servers. Use the report to identify servers that are running out of space.

Charts

The bar chart shows the capacity, used space, available space, and percentage of used space for disks on servers. Dots are used to represent the total space that is used by the file system, and the total space that is available to the file system on the server. You can position the cursor over the dot to show the amount of used space or the amount of available space for the file system. The position of the dot that is used to represent the total space that is available to the file system is also used to show the total capacity of the file system.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

For each server, the following information is provided:

Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Server Location

The physical location of a server. The location is defined when a server is added to IBM Spectrum® Control. You can add or edit the location of the server in the Properties pane of the server.

Server Total Disk Capacity (GiB)

The amount of storage space that is on a server.

Server Total Disk Used Space (GiB)

The amount of storage space that is used on all disks on a server.

Server Total Disk Available Space (GiB)

The amount of unused storage space on all disks on a server.

Server Total Disk Used Space (%)

The percentage of used storage space on all disks on the server.

Server Total File System Capacity (GiB)

The amount of storage space on the file system of the resource.

Server Total File System Used Space (GiB)

The amount of used storage space in the file system of the resource.

Server Total File System Available Space (GiB)

The amount of unused storage space in the file system of the resource.

Server Total File System Used Space (%)

The amount of used storage space in the file system of the resource.

Related reports

To open related reports, click the name of the server and then click the name of the report. You can open the following related reports:

- Server Disks Capacity
- File Systems Capacity

Related tasks

- [Running the Servers Capacity report](#)

Running the Most Active Servers report

To see which servers are the most active, run the Most Active Servers report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content, > IBM Spectrum Control Predefined Reports.
2. Click Servers.
3. Click Most Active Servers.
4. Select the servers, and then select a sort order.
5. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
6. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Servers report](#)

Most Active Servers report

Shows a chart of the 20 servers that have the most active volumes. The most active volumes have the greatest aggregate load.

Charts

The bar chart shows the most active servers. By default, servers are sorted by the total I/O rate.
Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active storage systems is as follows:

Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Volume Utilization

The average percentage of time that the volume is busy.

Related reports

To open the Summarized Performance of Volumes by Server report, click the name of the server in the report table.

Related tasks

- [Running the Most Active Servers report](#)

Running the Summarized Performance of Volumes by Server report

To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Server report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Servers.
3. Click Summarized Performance of Volumes by Server.
4. Select one or more servers.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Summarized Performance of Volumes by Server report](#)

Summarized Performance of Volumes by Server report

Shows summarized performance metrics for volumes on a server. Use the report to identify the aggregate load and the average response time of a particular server.

Charts

The charts show the following performance metrics for the volumes on a server:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, the write data rates, and the total data rates

- The read cache percentages and the write cache percentages

Report output

The following information is provided in the report for the period that you specified:

Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Day, hour, or time

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation, across all volumes on the server. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second, across all volumes on the server. This value includes both sequential and nonsequential operations. The average value for each volume is added to give a total I/O rate for all volumes on the server.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations, across all volumes on the server. The average value for each volume is added to give a total data rate for all volumes on the server.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation, across all volumes on the server.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation, across all volumes on the server.

Read Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations, across all volumes on the server. The average value for each volume is added to give a read data rate for all volumes on the server.

Write Data Rate (MiB/s)

The average number of MiB per second that are transferred for write operations, across all volumes on the server. The average value for each volume is added to give a write data rate for all volumes on the server.

Overall Read I/O Rate (ops/s)

The average number of read operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential write operations.

Overall Read Cache Hit Percentage

The average percentage of all read operations that find data in the cache, across all volumes on the server. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage

The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.

Related reports

To open the Performance of Volumes by Server report, click the name of the server in the report table.

Related tasks

- [Running the Summarized Performance of Volumes by Server report](#)

Running the Performance of Volumes by Server report

To see performance metrics for the volumes on a server, use the Performance of Volumes by Server report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Servers.
3. Click Performance of Volumes by Server.
4. Select a server.
5. Select the volumes.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of Volumes by Server report](#)

Performance of Volumes by Server report

Shows the performance metrics for volumes on a server. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

Charts

The charts show the following performance metrics for the volumes:

- The total overall I/O rates. The chart shows the contribution of each volume to the overall load on the server.
- The total data rates. The chart shows the contribution of each volume to the overall load on the server.
- The volume utilization percentages. The utilization percentage of each volume is displayed relative to the utilization percentage of the other volumes.
- The total response times. The response time of each volume is displayed relative to the response time of the other volumes.

Report output

The following information is provided in the report for the period that you specified:

Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

Day, hour, or time

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Volume Utilization

The average percentage of time that the volume is busy.

Related tasks

- [Running the Performance of Volumes by Server report](#)

Running the Server Disks Capacity report

To see information about the capacity of disks on servers, run the Server Disks Capacity report.

Before you begin

To run the report with the default settings, click Finish.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Servers.
3. Click Server Disks Capacity.
4. Optional: Select one or more servers.
5. Optional: Select a sort order.
6. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Server Disks Capacity report](#)

Server Disks Capacity report

Shows the available and used disk space for all disks, which are grouped by server. You can use the report to identify disks that have no allocated space.

Charts

The bar chart shows the capacity, used space, available space, and percentage of used space for disks on servers.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

For each server, the following information is provided:

- Server Disk Name
The identifier that specifies a path to a server disk, for example /dev/hdisk0.
- Server Disk Capacity (GiB)
The amount of storage space that is on a server disk.
- Server Disk Available Space (GiB)
The amount of unused storage space that is on a server disk.
- Server Disk Used Space (GiB)
The amount of storage space that is used on a server disk.
- Server Disk Used Space (%)
The percentage of used storage space on the disk.
- Server Disk Status
The condition of the resource, for example normal, warning, or error.
- Server Disk Firmware Version
The version number of the firmware that is running on the disk.
- Server Disk Model
The model name or model number of the resource.
- Server Disk Vendor
The vendor who supplied the resource.

Related tasks

- [Running the Server Disks Capacity report](#)

Predefined reports about file systems on servers

You can run predefined reports to analyze the capacity of file systems that are grouped by server.

- [Running the File Systems Capacity report](#)
To see information about the capacity of file systems on servers, run the File Systems Capacity report.
- [File Systems Capacity report](#)
Shows the available and used capacity of file systems, which are grouped by server. You can use this report to identify file systems that might be running out of space.

Running the File Systems Capacity report

To see information about the capacity of file systems on servers, run the File Systems Capacity report.

Before you begin

To run the report with the default settings, click Finish.

Tip: To select a resource, click the resource. To select multiple resources, press Ctrl and click the resources. To select a series of resources, select the first resource, and then press Shift and click the last resource.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Servers.
3. Click File Systems.
4. Click File Systems Capacity.
5. Select one or more servers.

6. Optional: Select a sort order.
7. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [File Systems Capacity report](#)

File Systems Capacity report

Shows the available and used capacity of file systems, which are grouped by server. You can use this report to identify file systems that might be running out of space.

Charts

The bar chart shows the capacity, used space, available space, and percentage of used space for file systems on servers.
Tip: To change the sort order, select a value from the Sort Order list.

Report output

For each server, the following information is provided:

File System Mount Point

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be c:\ or d:\. On operating systems such as AIX®, Linux®, Solaris, or HP-UX, the mount point might be /opt or /export/home.

File System Type

The type of file system that the resource uses.

File System Used Space (GiB)

The amount of used storage space in the file system of the resource.

File System Available Space (GiB)

The amount of unused storage space in the file system of the resource.

File System Capacity (GiB)

The amount of storage space on the file system of the resource.

File System Used Space (%)

The percentage of used space in the file system on the server.

File System Available Inodes

The number of unused inodes in file systems on the operating system.

File System Used Inodes

The number of used inodes in file systems on the operating system.

Tip: For Microsoft Windows systems, this property is blank.

Related tasks

- [Running the File Systems Capacity report](#)

Predefined reports about storage systems and components

You can run predefined reports to view capacity details and performance metrics for storage systems and their components. The components include controllers, modules, nodes, disks, I/O groups, managed disks, storage pools, ports, RAID arrays, and volumes.

- [Predefined reports about storage systems](#)
You can run predefined reports to view capacity details and performance metrics for storage systems.
- [Predefined reports about controllers, modules, and nodes](#)
You can run predefined reports to analyze and compare the performance of controllers, modules, and nodes.
- [Predefined reports about disks](#)
You can run predefined reports to analyze the capacity of disks and performance of local disks.
- [Predefined reports about host connections](#)
You can run predefined reports to analyze and compare the performance of volumes on host connections.
- [Predefined reports about I/O groups](#)
You can run predefined reports to analyze and compare the performance of I/O groups.
- [Predefined reports about managed disks](#)
You can run predefined reports about to analyze and compare the capacity and performance of managed disks.
- [Predefined reports about storage pools](#)
You can run predefined reports to analyze and compare the capacity and performance of storage pools.
- [Predefined reports about ports](#)
You can run predefined reports to analyze and compare the performance of ports.
- [Predefined reports about RAID arrays](#)
You can run predefined reports to analyze and compare the performance of RAID arrays.

- [Predefined reports about volumes](#)
You can run predefined reports to analyze the capacity and performance of volumes.

Predefined reports about storage systems

You can run predefined reports to view capacity details and performance metrics for storage systems.

- [Running the Storage Systems Capacity report](#)
To see information about the capacity of disks, pools, and volumes on storage systems, run the Storage Systems Capacity report.
- [Storage Systems Capacity report](#)
Shows a chart of the 20 storage systems with the most used space, and a table with details of all storage systems. Use the report to review the space allocation in your storage systems.
- [Running the Storage Systems Historical Capacity report](#)
To see changes in the usage of space on storage systems over a specified period, run the Storage Systems Historical Capacity report.
- [Storage Systems Historical Capacity report](#)
Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.
- [Running the Most Active Storage Systems report](#)
To see which storage systems are the most active, run the Most Active Storage Systems report.
- [Most Active Storage Systems report](#)
Shows a chart of the 20 storage systems that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of storage systems.
- [Running the Performance of One Storage System report](#)
To see performance metrics for a storage system, run the Performance of One Storage System report.
- [Performance of One Storage System report](#)
Shows four charts, and a table of performance metrics for a storage system over a time period that you specify. Use the report to analyze the performance of a storage system.
- [Running the Compare Performance of Multiple Storage Systems report](#)
Use the Compare Performance of Multiple Storage Systems report to compare up to four performance metrics on multiple storage systems. For example, you can compare cache percentages, data rates, and I/O rates.
- [Compare Performance of Multiple Storage Systems report](#)
Shows up to four performance metrics for multiple storage systems over time. Use the report to compare the performance metrics for storage systems.
- [Running the Compare Performance of One Storage System over Time Ranges report](#)
To compare a performance metric for one storage system in two date ranges, run the Compare Performance of One Storage System over Time Ranges report.
- [Compare Performance of One Storage System over Time Ranges report](#)
Shows one performance metric on one storage system over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.
- [Running the Performance Data Export report](#)
To get a spreadsheet of the performance metrics for one or more resources on a storage system, run the Performance Data Export report.
- [Performance Data Export report](#)
Exports the performance metrics for one or more resources on a storage system to a spreadsheet file in Microsoft Excel file format. Use the report if you want to analyze performance data in Excel or to send performance data to IBM®.


Running the Storage Systems Capacity report

To see information about the capacity of disks, pools, and volumes on storage systems, run the Storage Systems Capacity report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content  IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Storage Systems Capacity.
4. Optional: Select the configuration of the storage systems.
5. Select one or more storage systems.
6. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Storage Systems Capacity report](#)

Storage Systems Capacity report

Shows a chart of the 20 storage systems with the most used space, and a table with details of all storage systems. Use the report to review the space allocation in your storage systems.

Charts

The bar chart shows the used space and available space for pools on storage systems.

Report output

Note: IBM FlashSystem® A9000 and A9000R storage systems are not supported by IBM Spectrum® Control for this report.
For each storage system, the following information is provided:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage System Type

The type of storage system. For example, the storage system can be an IBM System Storage DS8800 system, an IBM System Storage DS8700 system, an IBM System Storage XIV system, or another type of storage system.

Storage System Location

The physical location of a storage system. The location is defined when a storage system is added to IBM Spectrum Control. You can add or edit the location of the storage system in the Properties pane of the storage system.

Storage System Pool Capacity (GiB)

The amount of storage space in pools that are on the storage system. For an XIV® or IBM Spectrum Accelerate, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage System Used Pool Space (GiB)

The amount of space that is in use in all pools on a storage system.

Storage System Pool Available Space (GiB)

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

Storage System Physical Allocation Percentage

The percentage of physical space in storage system pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{allocated pool space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a 200 GiB storage pool. Therefore, the space that is reserved for volumes is 50 GiB.

Storage System Virtual Allocation Percentage

The percentage of physical space in storage system pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage System Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

Storage System Shortfall Percentage

The percentage of the remaining unallocated volume space in storage system pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{volume space} - \text{used volume space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage System Volume Capacity (GiB)

The amount of space on all volumes on the storage system.

Storage System Number of Volumes

The number of volumes on a resource.

Storage System Number of Disks

The number of physical disks on the storage system. For a resource that is running IBM Spectrum Virtualize and is configured as a back-end device, the value is the number of managed disks on the resource.

Storage System Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage System Unallocated Disk Space (GiB)

The amount of disk space that can be added to a pool.

Storage System Volume Capacity Assigned to MDisks (GiB)

The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

Storage System Disk Capacity (GiB)

The total amount of space that is on a physical disk on the storage system.

Storage System Total Disk Capacity (GiB)

The amount of space on physical disks on a storage system, including spare disks.

Storage System Cache (GiB)

The size of the cache on the storage system. This value is not shown for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage System Assigned Volume Space (GiB)

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage System Volume Capacity for z/OS

The amount of space on all volumes on the storage system that the z/OS® operating system can use.

Storage System Real Available Pool Space (GiB)

The amount of unused space in pools that have an associated soft size on an XIV or IBM Spectrum Accelerate. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Real Configured Pool Space (GiB)

The amount of storage space that is in pools that have an associated soft size on an XIV or IBM Spectrum Accelerate. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System MDisk Available Space (GiB)

The amount of storage space that is available on a managed disk. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Storage Systems Historical Capacity
- Disks Capacity
- Volumes Capacity
- Pools Capacity

Related tasks

- [Running the Storage Systems Capacity report](#)

Running the Storage Systems Historical Capacity report

To see changes in the usage of space on storage systems over a specified period, run the Storage Systems Historical Capacity report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Storage Systems Historical Capacity.
4. Optional: Select the configuration of the storage systems.
5. Select one or more storage systems.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Storage Systems Historical Capacity report](#)

Storage Systems Historical Capacity report

Shows the usage of space on storage systems over a specified period. Use the report to analyze changes in your storage system requirements.

Charts

The chart shows the following space statistics for the storage systems:

- The capacity of the pools on the storage systems
- The used space in the pools on the storage systems
- The available space in the pools on the storage systems
- The capacity of the volumes on the storage systems
- The available space on the managed disks on the storage systems
- The capacity of the managed disks on the storage systems

Report output

The following information is included in reports for storage systems:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Date

The date that the data was collected for the report.

Storage System Pool Capacity (GiB)

The amount of storage space in pools that are on the storage system. For an XIV® or IBM Spectrum Accelerate, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage System Used Pool Space (GiB)

The amount of space that is in use in all pools on a storage system.

Storage System Available Pool Space (GiB)

The amount of unused space in pools that are on the storage system.

Storage System Volume Capacity (GiB)

The amount of space on all volumes on the storage system.

Storage System Number of Volumes

The number of volumes on a resource.

Storage System Number of Disks

The number of physical disks on the storage system. For a resource that is running IBM Spectrum Virtualize and is configured as a back-end device, the value is the number of managed disks on the resource.

Storage System Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage System Unallocated Disk Space (GiB)

The amount of disk space that can be added to a pool.

Storage System Volume Capacity Assigned to MDisks (GiB)

The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

Storage System Disk Capacity (GiB)

The amount of storage space on physical disks in a storage system, excluding spare disks. This value is not shown for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage System Total Disk Capacity (GiB)

The amount of space on physical disks on a storage system, including spare disks.

Storage System Assigned Volume Space (GiB)

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage System Volume Capacity for z/OS (GiB)

The amount of space on all volumes on the storage system that the z/OS® operating system can use.

Storage System Real Available Pool Space (GiB)

The amount of unused space in pools that have an associated soft size on an XIV or IBM Spectrum Accelerate. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Real Configured Pool Space (GiB)

The amount of storage space that is in pools that have an associated soft size on an XIV or IBM Spectrum Accelerate. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System MDisk Capacity (GiB)

The amount of storage space on the managed disk.

Storage System MDisk Available Space (GiB)

The amount of storage space that is available on a managed disk. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Related tasks

- [Running the Storage Systems Historical Capacity report](#)

Running the Most Active Storage Systems report

To see which storage systems are the most active, run the Most Active Storage Systems report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Most Active Storage Systems.
4. Optional: Select the configuration of the storage systems.
5. Select the storage systems, and then select a sort order.
6. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
7. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Storage Systems report](#)

Most Active Storage Systems report

Shows a chart of the 20 storage systems that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of storage systems.

Charts

The bar chart shows the most active storage systems. By default, storage systems are sorted by the total I/O rate.
Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active storage systems is as follows:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Location

The physical location of a storage system. The location is defined when a storage system is added to IBM Spectrum Control. You can add or edit the location of the storage system in the Properties pane of the storage system.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Most Active Host Connections
- Most Active Nodes
- Most Active IO Groups
- Most Active Disks
- Most Active Managed Disks
- Most Active Pools
- Most Active Ports
- Most Active Volumes
- Most Active Controllers or Modules
- Most Active RAID Arrays
- Performance of One Storage System
- Storage Resource Relationships Summary

Related tasks

- [Running the Most Active Storage Systems report](#)

Running the Performance of One Storage System report

To see performance metrics for a storage system, run the Performance of One Storage System report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Performance of One Storage System.
4. Optional: Select the configuration of the storage system.
5. Select a storage system.
6. Optional: Select an interval.
7. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One Storage System report](#)

Performance of One Storage System report

Shows four charts, and a table of performance metrics for a storage system over a time period that you specify. Use the report to analyze the performance of a storage system.

Charts

The charts show the following performance metrics for a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, write I/O rates, and response times
- The read data rates, and the write data rates
- The read cache hit percentages, and the write cache hit percentages

Report output

The following information is provided in the report for the period that you specified:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Location

The physical location of a storage system. The location is defined when a storage system is added to IBM Spectrum Control. You can add or edit the location of the storage system in the Properties pane of the storage system.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Most Active Host Connections
- Most Active Nodes
- Most Active IO Groups
- Most Active Disks
- Most Active Managed Disks
- Most Active Pools
- Most Active Ports
- Most Active Volumes
- Most Active Controllers or Modules
- Most Active RAID Arrays

Related tasks

- [Running the Performance of One Storage System report](#)

Running the Compare Performance of Multiple Storage Systems report

Use the Compare Performance of Multiple Storage Systems report to compare up to four performance metrics on multiple storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Compare Performance of Multiple Storage Systems.
4. Select the storage systems, and then click Next.
5. Select the category of performance metric that you want to include in the report, and then select the performance metric.
6. Optional: To add more metrics to the report, complete the following steps:
 - a. Click Add Metric.
 - b. Select the category of performance metric, and then select the performance metric.
 - c. If you want to remove the last metric that you added, click Remove Metric.
7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period.

It is easier to analyze small amounts of information over shorter periods.

Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
9. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple Storage Systems report](#)

Compare Performance of Multiple Storage Systems report

Shows up to four performance metrics for multiple storage systems over time. Use the report to compare the performance metrics for storage systems.

Charts

The charts show the performance metrics for the storage systems that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metrics

Shows the data for each performance metric.

Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

Related tasks

- [Running the Compare Performance of Multiple Storage Systems report](#)

Running the Compare Performance of One Storage System over Time Ranges report

To compare a performance metric for one storage system in two date ranges, run the Compare Performance of One Storage System over Time Ranges report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Compare Performance of One Storage System over Time Ranges.
4. Optional: Select the configuration of the storage systems.
5. Select the storage system.
6. Select the category of performance metric about which you want a report, and then select the performance metric.
7. Select an interval.
8. Specify the two date ranges or reporting periods that you want to compare.
For example, you can compare the performance metric for this week with last week.
Tip: To specify a start date and an end date for a reporting period, select Custom Date Range from the list of reporting periods.
9. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Compare Performance of One Storage System over Time Ranges report](#)

Compare Performance of One Storage System over Time Ranges report

Shows one performance metric on one storage system over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Charts

The chart shows the performance data of the storage system on the performance metric you selected, over the time periods that you selected.

Report output

The following information is provided in the report:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Day in Range

The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The number of days that are shown is determined by the reporting period that was selected for the report. If Last 30 Days was selected, then entries for days 1 - 30 are shown on the report.

Hour in Range

The number of the hour in the day at which the performance metric data was collected. For example, 0 is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

Minute in Range

The interval of minutes during which the performance metric data was collected. For example, 0 shows that performance data was collected 0 - 5 minutes after the hour. Similarly, 5 shows that the performance data that was collected 5 - 10 minutes after the hour.

First Date Range

The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the first date range.

Second Date Range

The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the second date range.

Related tasks

- [Running the Compare Performance of One Storage System over Time Ranges report](#)

Running the Performance Data Export report

To get a spreadsheet of the performance metrics for one or more resources on a storage system, run the Performance Data Export report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Performance Data Export.
4. Optional: Select the configuration of the storage systems.
5. Select one or more storage systems.
6. Optional: Click Next, and then select the devices for which you want performance metrics in the report.
7. Click Finish.

Results

You can open or save the spreadsheet from your browser.

Related reference

- [Performance Data Export report](#)

Performance Data Export report

Exports the performance metrics for one or more resources on a storage system to a spreadsheet file in Microsoft Excel file format. Use the report if you want to analyze performance data in Excel or to send performance data to IBM®.

Report output

The spreadsheet contains a worksheet for each resource on the storage system. The worksheet for a resource contains the performance metric data for that resource.

Related tasks

- [Running the Performance Data Export report](#)

Predefined reports about controllers, modules, and nodes

You can run predefined reports to analyze and compare the performance of controllers, modules, and nodes.

- [Running the Most Active Controllers or Modules report](#)
To see which controllers or modules are most active on storage systems, run the Most Active Controllers or Modules report.
- [Most Active Controllers or Modules report](#)
Shows a chart of the 20 controllers or modules that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of controllers or modules.
- [Running the Performance of One Controller or Module report](#)
To see performance metrics for controllers or modules on a storage system, run the Performance of One Controller or Module report.
- [Performance of One Controller or Module report](#)
Shows four charts, and a table of performance metrics for a controller or module over a time period. Use the report to analyze the performance of a controller or module.
- [Running the Most Active Nodes report](#)
To see which nodes are most active on your storage systems, run the Most Active Nodes report.
- [Most Active Nodes report](#)
Shows a chart of the 20 nodes that are most active, and details for all nodes, for a time period that you specify. Use the report to analyze the performance of nodes.
- [Running the Performance of One Node report](#)
To see information about the performance of a node on a storage system, run the Performance of One Node report.
- [Performance of One Node report](#)
Shows four charts, and a table of performance metrics for a node over a time period. Use the report to analyze the performance of a node.
- [Running the Compare Performance of Multiple Controllers and Modules report](#)
Use this report to compare up to four performance metrics for multiple controllers and modules that are on storage systems. For example, you can compare data

rates, I/O rates, and response times.

- [Compare Performance of Multiple Controllers and Modules report](#)

Shows up to four performance metrics for multiple controllers or modules over time. Use the report to compare the performance metrics for controllers or modules that are on storage systems.

- [Running the Compare Performance of Multiple Nodes report](#)

Use the Compare Performance of Multiple Controllers and Modules report to compare up to four performance metrics on multiple nodes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

- [Compare Performance of Multiple Nodes report](#)

Shows up to four performance metrics for multiple nodes over time. Use the report to compare the performance metrics for nodes that are on storage systems.

Running the Most Active Controllers or Modules report

To see which controllers or modules are most active on storage systems, run the Most Active Controllers or Modules report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Controllers, Modules, and Nodes.
4. Click Most Active Controllers or Modules.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Controllers or Modules report](#)

Most Active Controllers or Modules report

Shows a chart of the 20 controllers or modules that are most active, and details for all storage systems, for a time period that you specify. Use the report to analyze the performance of controllers or modules.

Restriction: This report applies only to DS8000®, XIV® systems, IBM Spectrum Accelerate, and vendor-acquired systems.

Charts

The bar chart shows the most active controllers or modules on storage systems. By default, controllers or modules are sorted by the overall I/O rate.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active controllers or modules is as follows:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Controller or Module Name

The name that was assigned to the controller or module when it was added to the system.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Controller or module	Performance of One Controller or Module

Related tasks

- [Running the Most Active Controllers or Modules report](#)

Running the Performance of One Controller or Module report

To see performance metrics for controllers or modules on a storage system, run the Performance of One Controller or Module report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Controllers, Modules, and Nodes.
4. Click Performance of One Controller or Module.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a controller or module.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
10. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One Controller or Module report](#)

Performance of One Controller or Module report

Shows four charts, and a table of performance metrics for a controller or module over a time period. Use the report to analyze the performance of a controller or module.

Restriction: This report applies only to DS8000®, XIV® systems, IBM Spectrum Accelerate, and vendor-acquired systems.

Charts

The charts show the following performance metrics for each controller or module on a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, and the response times
- The read data rates, and the write data rates
- The overall cache hit percentage for read operations and write operations, and the utilization percentage for processors

Report output

The following information is provided in the report for the period that you specified:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Controller Name or Module Name

The name that was assigned to the controller or module when it was added to the system.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Related reports

To open the Most Active Controllers or Modules report, click the name of the controller or module in the report table.

Related tasks

- [Running the Performance of One Controller or Module report](#)

Running the Most Active Nodes report

To see which nodes are most active on your storage systems, run the Most Active Nodes report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Controllers, Modules, and Nodes.
4. Click Most Active Nodes.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Nodes report](#)

Most Active Nodes report

Shows a chart of the 20 nodes that are most active, and details for all nodes, for a time period that you specify. Use the report to analyze the performance of nodes.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Charts

The chart shows the most active nodes on storage systems. By default, nodes are sorted by the total overall I/O rate.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided for nodes is as follows:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage I/O Group Name

The name that was assigned to the I/O group when it was added to the system.

Storage Node Name

The name that was assigned to the storage system node when it was added to the system.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Processor Utilization Percentage

The average percentage of time that the processors on nodes are busy doing system I/O tasks. This value applies only to resources that are running IBM Spectrum Virtualize.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
I/O group	Performance of One IO Group
Node	Performance of One Node

Related tasks

- [Running the Most Active Nodes report](#)

Running the Performance of One Node report

To see information about the performance of a node on a storage system, run the Performance of One Node report.

Procedure

1. In the Welcome portal, click Team Content, IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Controllers, Modules, and Nodes.
4. Click Performance of One Node.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a node.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
10. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One Node report](#)

Performance of One Node report

Shows four charts, and a table of performance metrics for a node over a time period. Use the report to analyze the performance of a node.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Charts

The charts show the following performance statistics for a node on a storage system:

- The total overall I/O rates, and the total response times
- The overall I/O rates for read and write operations, the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

Report output

The following information is provided in the report for the period that you specified:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage I/O Group Name

The name that was assigned to the I/O group when it was added to the system.

Storage Node Name

The name that was assigned to the storage system node when it was added to the system.

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Write Cache Delay Percentage

The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Processor Utilization Percentage

The average percentage of time that the processors on nodes are busy doing system I/O tasks. This value applies only to resources that are running IBM Spectrum Virtualize.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Most Active Nodes
I/O group	Performance of One IO Group

Related tasks

- [Running the Performance of One Node report](#)

Running the Compare Performance of Multiple Controllers and Modules report

Use this report to compare up to four performance metrics for multiple controllers and modules that are on storage systems. For example, you can compare data rates, I/O rates, and response times.

Procedure

1. In the Welcome portal, click Team Content, IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Controllers, Modules, and Nodes.
4. Click Compare Performance of Multiple Controllers and Modules.
5. Select the storage systems, and then click Refresh.
6. Select the controllers and modules, and then click Next.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
 - a. Click Add Metric.
 - b. Select the category of performance metric, and then select the performance metric.
 - c. If you want to remove the last metric that you added, click Remove Metric.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period.

It is easier to analyze small amounts of information over shorter periods.

Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
11. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple Controllers and Modules report](#)

Compare Performance of Multiple Controllers and Modules report

Shows up to four performance metrics for multiple controllers or modules over time. Use the report to compare the performance metrics for controllers or modules that are on storage systems.

Restriction: This report applies only to DS8000®, XIV® systems, IBM Spectrum Accelerate, and vendor-acquired systems.

Charts

The charts show the performance metrics for the controllers and modules that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metric

Shows the time at which the performance metric data was collected.

Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

Controller and modules

Shows the names of the controllers and modules, and the performance metric data for each controller and module.

Related tasks

- [Running the Compare Performance of Multiple Controllers and Modules report](#)

Running the Compare Performance of Multiple Nodes report

Use the Compare Performance of Multiple Controllers and Modules report to compare up to four performance metrics on multiple nodes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Controllers, Modules, and Nodes.
4. Click Compare Performance of Multiple Nodes.
5. Select the storage systems, and then click Refresh.
6. Select the nodes, and then click Next.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
 - a. Click Add Metric.
 - b. Select the category of performance metric, and then select the performance metric.
 - c. If you want to remove the last metric that you added, click Remove Metric.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
11. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple Nodes report](#)

Compare Performance of Multiple Nodes report

Shows up to four performance metrics for multiple nodes over time. Use the report to compare the performance metrics for nodes that are on storage systems.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Charts

The charts show the performance metrics for the nodes that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metrics

Shows the data for each performance metric.

Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

Nodes

Shows the names of the nodes, and the performance metric data for each node.

Related tasks

- [Running the Compare Performance of Multiple Nodes report](#)

Predefined reports about disks

You can run predefined reports to analyze the capacity of disks and performance of local disks.

- [Running the Disks Capacity report](#)
Run the Disks Capacity report to see information about the capacity of disks on storage systems.
- [Disks Capacity report](#)
Shows a chart of the disk classes with the most allocated space, and a table with space statistics for all disks by class. Use the report to review space allocation in your disks.
- [Running the Most Active Disks report](#)
To see which local disks that are on storage systems are the most active, use the Most Active Disks report. This report shows information only for SAN Volume Controller and Storwize® V7000 storage systems.
- [Most Active Disks report](#)
Shows a chart of the 20 local disks that are most active, and details for all local disks, for a time period that you specify. Use the report to analyze the performance of local disks.
- [Running the Performance of One Disk report](#)
To see performance metrics for a local disk on a storage system, use the Performance of One Disk report.

- [Performance of One Disk report](#)
Shows four charts, and a table of performance metrics for a local disk over a time period. Use the report to analyze the performance of a local disk.
- [Running the Compare Performance of Multiple Disks report](#)
Use the Compare Performance of Multiple Disks report to compare up to four performance metrics for multiple local disks that are on storage systems. You can view reports about local disks on SAN Volume Controller, Storwize V7000, Storwize V7000 Unified, and Storwize V3700 systems. For example, you can compare cache percentages, data rates, and I/O rates.
- [Compare Performance of Multiple Disks report](#)
Shows up to four performance metrics for multiple local disks over time. Use the report to compare the performance metrics for local disks that are on storage systems.

Running the Disks Capacity report

Run the Disks Capacity report to see information about the capacity of disks on storage systems.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Disks.
4. Click Disks Capacity.
5. Optional: Select the configuration of the storage systems.
6. Select one or more storage systems.
7. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Disks Capacity report](#)

Disks Capacity report

Shows a chart of the disk classes with the most allocated space, and a table with space statistics for all disks by class. Use the report to review space allocation in your disks.

Charts

The bar chart shows the capacity of disks, grouped by storage class, on storage systems.

Report output

For each disk on a storage system, the following information is provided:

- Storage Disk Name
The name that was assigned to the storage disk when it was added to the system.
- Storage System Type
The type of storage system. For example, the storage system can be an IBM System Storage DS8800 system, an IBM System Storage DS8700 system, an IBM System Storage XIV system, or another type of storage system.
- Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.
- Storage Pool Name
The name that was assigned to the pool when it was added to the system.
- Storage System Name
A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.
- Storage Disk Tag
A number that identifies the array to which a disk belongs. This property applies only to DS8000® storage systems.
- Storage Disk Vendor
The vendor who supplied the resource.
- Storage Disk Model
The model name or model number of the resource.
- Storage Disk Serial Number

The serial number of the resource.

Storage Disk Firmware Version
The version number of the firmware that is running on the disk.

Storage Disk Capacity (GiB)
The amount of storage space that is on a physical disk on the storage system.

Storage Disk Speed (RPM)
The speed of the disk.

Related tasks

- [Running the Disks Capacity report](#)

Running the Most Active Disks report

To see which local disks that are on storage systems are the most active, use the Most Active Disks report. This report shows information only for SAN Volume Controller and Storwize® V7000 storage systems.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Disks.
4. Click Most Active Disks.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Disks report](#)

Most Active Disks report

Shows a chart of the 20 local disks that are most active, and details for all local disks, for a time period that you specify. Use the report to analyze the performance of local disks.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Charts

The bar chart shows the most active local disks that are on storage systems. By default, local disks are sorted by the total I/O rate.
Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active local disks that are on storage systems is as follows:

Storage System Name
A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage MDisk Name
The name that was assigned to the managed disk when it was added to the system.

Storage Local Disk Name
The name that was assigned to the local disk on the storage system when the local disk was created on the system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage Local Disk Storage Class

The storage technology of the local disk on a storage system. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD).

Storage Local Disk Speed (RPM)

The speed of the local disk on a storage system.

Storage Local Disk Capacity (GiB)

The amount of storage space on the local disk.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Managed disk	Performance of One Managed Disk
Disk	Performance of One Disk

Related tasks

- [Running the Most Active Disks report](#)

Running the Performance of One Disk report

To see performance metrics for a local disk on a storage system, use the Performance of One Disk report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Disks.
4. Click Performance of One Disk.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a local disk.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
10. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One Disk report](#)

Performance of One Disk report

Shows four charts, and a table of performance metrics for a local disk over a time period. Use the report to analyze the performance of a local disk.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Charts

The charts show the following performance metrics for a local disk on a storage system:

- The total back-end I/O rates and the overall back-end response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, the write data rates, and the total data rates

- The read queue times, the write queue times, and the overall queue times

Report output

The following information is provided in the report for the period that you specified:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage MDisk Name

The name that was assigned to the managed disk when it was added to the system.

Storage Local Disk Name

The name that was assigned to the local disk on the storage system when the local disk was created on the system.

Storage Local Disk Storage Class

The storage technology of the local disk on a storage system. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD).

Storage Local Disk Speed (RPM)

The speed of the local disk on a storage system.

Storage Local Disk Capacity (GiB)

The amount of storage space on the local disk.

Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)

The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Queue Time (ms/op)

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Related reports

To open the Most Active Disks report, click the name of the storage system in the report table.

Related tasks

- [Running the Performance of One Disk report](#)

Running the Compare Performance of Multiple Disks report

Use the Compare Performance of Multiple Disks report to compare up to four performance metrics for multiple local disks that are on storage systems. You can view reports about local disks on SAN Volume Controller, Storwize® V7000, Storwize V7000 Unified, and Storwize V3700 systems. For example, you can compare cache percentages, data rates, and I/O rates.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Disks.
4. Click Compare Performance of Multiple Disks.
5. Select the storage systems, and then click Refresh.
6. Select the local disks, and then click Next.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
 - a. Click Add Metric.

- b. Select the category of performance metric, and then select the performance metric.
 - c. If you want to remove the last metric that you added, click Remove Metric.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
11. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple Disks report](#)

Compare Performance of Multiple Disks report

Shows up to four performance metrics for multiple local disks over time. Use the report to compare the performance metrics for local disks that are on storage systems.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Charts

The charts show the performance metrics for the local disks that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metric

Shows the time at which the performance metric data was collected.

Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

Disks

Shows the names of the disks, and the performance metric data for each disk.

Related tasks

- [Running the Compare Performance of Multiple Disks report](#)

Predefined reports about host connections

You can run predefined reports to analyze and compare the performance of volumes on host connections.

- [Running the Most Active Host Connections report](#)
To see which host connections are the most active, run the Most Active Host Connections report.
- [Most Active Host Connections report](#)
Shows a chart of the 20 host connections that have the most active volumes. The most active volumes have the greatest aggregate load.
- [Running the Summarized Performance of Volumes by Host Connection report](#)
To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Host Connection report.
- [Summarized Performance of Volumes by Host Connection report](#)
Shows summarized performance metrics for volumes on a host connection. Use the report to identify the aggregate load and the average response time of a particular server.
- [Running the Performance of Volumes by Host Connection report](#)
To see performance metrics for the volumes on a server, use the Performance of Volumes by Host Connection report.
- [Performance of Volumes by Host Connection report](#)
Shows the performance metrics for volumes on a host connection. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

Running the Most Active Host Connections report

To see which host connections are the most active, run the Most Active Host Connections report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Host Connections.
4. Click Most Active Host Connections.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Host Connections report](#)

Most Active Host Connections report

Shows a chart of the 20 host connections that have the most active volumes. The most active volumes have the greatest aggregate load.

Chart

The bar chart shows the most active host connections. By default, host connections are sorted by the total I/O rate.
Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active host connections is as follows:

Storage Host Connection Name

The user-defined name that describes the server or storage device that is assigned to a volume.

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

XIV® Cluster

The name of the cluster, if any, as defined on the XIV. If no cluster is defined for the host connection, this field is blank.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Host connection	Summarized Performance of Volumes by Host Connection
Storage system	Performance of One Storage System

Related tasks

- [Running the Most Active Host Connections report](#)

Running the Summarized Performance of Volumes by Host Connection report

To see the load that a particular server adds to your storage environment, use the Summarized Performance of Volumes by Host Connection report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Host Connections.
4. Click Summarized Performance of Volumes by Host Connection.
5. Select a host connection. To select a host connection, complete the following steps:
 - a. Type a keyword.
For example, type `example1` to find the names of host connections that start with `example1`.
To change the default search option, click Options.
Use the percent sign as a wildcard.
For example, type `a%` to find host connections that begin with 'a' or 'A'. Type `%a` to find host connections that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the host connections on a resource.
 - b. Click Search.
The report searches the user-defined names for the host connections. The report does not search the server names that are defined when a server is probed.
If the search is successful, the host connections are displayed in the Choice list.
 - c. From the list, select a host connection, and then click Refresh.
6. Select the volumes.
7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
9. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Summarized Performance of Volumes by Host Connection report](#)

Summarized Performance of Volumes by Host Connection report

Shows summarized performance metrics for volumes on a host connection. Use the report to identify the aggregate load and the average response time of a particular server.

Charts

The charts show the following performance statistics for the volumes on a host connection:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, the write data rates, and the total data rates
- The read cache percentages and the write cache percentages

Report output

The following information is provided in the report for the period that you specified:

Storage Host Connection Name

The user-defined name that describes the server or storage device that is assigned to a volume.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation, across all volumes on the server. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second, across all volumes on the server. This value includes both sequential and nonsequential operations. The average value for each volume is added to give a total I/O rate for all volumes on the server.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations, across all volumes on the server. The average value for each volume is added to give a total data rate for all volumes on the server.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation, across all volumes on the server.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation, across all volumes on the server.

Read Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations, across all volumes on the server. The average value for each volume is added to give a read data rate for all volumes on the server.

Write Data Rate (MiB/s)

The average number of MiB per second that are transferred for write operations, across all volumes on the server. The average value for each volume is added to give a write data rate for all volumes on the server.

Overall Read I/O Rate (ops/s)

The average number of read operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second, which is totaled for all volumes on the server. This value includes both sequential and nonsequential write operations.

Overall Read Cache Hit Percentage

The average percentage of all read operations that find data in the cache, across all volumes on the server. This value includes both sequential and nonsequential read operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Write Cache Hit Percentage

The average percentage of all write operations that are handled in the cache, across all volumes on the server. This value includes both sequential and nonsequential write operations.

Related reports

To open the Performance of Volumes by Host Connection report, click the name of the host connection in the report table.

Related tasks

- [Running the Summarized Performance of Volumes by Host Connection report](#)

Running the Performance of Volumes by Host Connection report

To see performance metrics for the volumes on a server, use the Performance of Volumes by Host Connection report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Host Connections.
4. Click Performance of Volumes by Host Connection.
5. Select a host connection. To select a host connection, complete the following steps:
 - a. Type a keyword.
For example, type `example1` to find the names of host connections that start with `example1`.
To change the default search option, click Options.
Use the percent sign as a wildcard.
For example, type `a%` to find host connections that begin with 'a' or 'A'. Type `%a` to find host connections that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the host connections on a resource.
 - b. Click Search.
The report searches the user-defined names for the host connections. The report does not search the server names that are defined when a server is probed.
If the search is successful, the host connections are displayed in the Choice list.
 - c. From the list, select a host connection, and then click Refresh.
6. Select the volumes.
7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
9. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of Volumes by Host Connection report](#)

Performance of Volumes by Host Connection report

Shows the performance metrics for volumes on a host connection. Use the report to identify the contribution of individual volumes to the load of a particular server in your storage environment.

Charts

The charts show the following performance metrics for the volumes:

- The total overall I/O rates. This is a stacked area chart that shows the contribution of each volume to the overall load on the server.
- The total data rates. This is a stacked area chart that shows the contribution of each volume to the overall load on the server.
- The volume utilization percentages. The utilization percentage of each volume is displayed relative to the utilization percentage of the other volumes.
- The total response times. The response time of each volume is displayed relative to the response time of the other volumes.

Report output

The following information is provided in the report for the period that you specified:

Storage Host Connection Name

The user-defined name that describes the server or storage device that is assigned to a volume.

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Volume Utilization

The average percentage of time that the volume is busy.

Related tasks

- [Running the Performance of Volumes by Host Connection report](#)

Predefined reports about I/O groups

You can run predefined reports to analyze and compare the performance of I/O groups.

- [Running the Most Active IO Groups report](#)
To see which I/O groups are most active on storage systems, run the Most Active IO Groups report.
- [Most Active IO Groups report](#)
Shows a chart of the 20 I/O groups that are most active, and details for all I/O groups, for a time period that you specify. Use the report to analyze the performance of I/O groups.
- [Running the Performance of One IO Group report](#)
To see information about the performance of an I/O group on a storage system, run the Performance of One IO Group report.
- [Performance of One IO Group report](#)
Shows four charts, and a table of performance metrics for an I/O group over a time period. Use the report to analyze the performance of an I/O group.
- [Running the Compare Performance of Multiple IO Groups report](#)
Use the Compare Performance of Multiple IO Groups report to compare up to four performance metrics for multiple I/O groups that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.
- [Compare Performance of Multiple IO Groups report](#)
Shows up to four performance metrics for multiple I/O groups over time. Use the report to compare the performance metrics for I/O groups that are on storage systems.

Running the Most Active IO Groups report

To see which I/O groups are most active on storage systems, run the Most Active IO Groups report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content, IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click IO Groups.
4. Click Most Active IO Groups.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.

8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active IO Groups report](#)

Most Active IO Groups report

Shows a chart of the 20 I/O groups that are most active, and details for all I/O groups, for a time period that you specify. Use the report to analyze the performance of I/O groups.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Charts

The chart shows the most active I/O groups on storage systems. By default, I/O groups are sorted by the total overall I/O rate.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided for I/O groups is as follows:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage I/O Group Name

The name that was assigned to the I/O group when it was added to the system.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Processor Utilization Percentage

The average percentage of time that the processors on nodes are busy doing system I/O tasks. This value applies only to resources that are running IBM Spectrum Virtualize.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
I/O group	Performance of One IO Group

Related tasks

- [Running the Most Active IO Groups report](#)

Running the Performance of One IO Group report

To see information about the performance of an I/O group on a storage system, run the Performance of One IO Group report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click IO Groups.
4. Click Performance of One IO Group.
5. Optional: Select the configuration of the storage systems.

6. Select a storage system.
7. Select an I/O group.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
10. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One IO Group report](#)

Performance of One IO Group report

Shows four charts, and a table of performance metrics for an I/O group over a time period. Use the report to analyze the performance of an I/O group.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Charts

The charts show the following performance metrics for an I/O group on a storage system:

- The total overall I/O rates, and the total response times
- The overall I/O rates for read and write operations, and the response times for read and write operations
- The data rates for read and write operations, and the total data rates
- The overall cache hit percentage for read operations, the cache delay percentage for write operations, and the utilization percentage for processors

Report output

The following information is provided in the report for the period that you specified:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage I/O Group Name

The name that was assigned to the I/O group when it was added to the system.

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Write Cache Delay Percentage

The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Processor Utilization Percentage

The average percentage of time that the processors on nodes are busy doing system I/O tasks. This value applies only to resources that are running IBM Spectrum Virtualize.

Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Performance of One Storage System
- Most Active IO Groups

Related tasks

- [Running the Performance of One IO Group report](#)

Running the Compare Performance of Multiple IO Groups report

Use the Compare Performance of Multiple IO Groups report to compare up to four performance metrics for multiple I/O groups that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click IO Groups.
4. Click Compare Performance of Multiple IO Groups.
5. Select the storage systems, and then click Refresh.
6. Select the I/O groups, and then click Next.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
 - a. Click Add Metric.
 - b. Select the category of performance metric, and then select the performance metric.
 - c. If you want to remove the last metric that you added, click Remove Metric.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period.

It is easier to analyze small amounts of information over shorter periods.

Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
11. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple IO Groups report](#)

Compare Performance of Multiple IO Groups report

Shows up to four performance metrics for multiple I/O groups over time. Use the report to compare the performance metrics for I/O groups that are on storage systems.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Charts

The charts show the performance metrics for the I/O groups that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metric

Shows the time at which the performance metric data was collected.

Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

I/O groups

Shows the names of the I/O groups, and the performance metric data for each I/O group.

Related tasks

- [Running the Compare Performance of Multiple IO Groups report](#)

Predefined reports about managed disks

You can run predefined reports about to analyze and compare the capacity and performance of managed disks.

- [Running the Managed Disks Capacity report](#)
To see information about used space and available space on managed disks (MDisks) on storage virtualizers, run the Managed Disks Capacity report.
- [Managed Disks Capacity report](#)
Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.
- [Running the Most Active Managed Disks report](#)
To see which managed disks are most active on storage systems, use the Most Active Managed Disks.
- [Most Active Managed Disks report](#)
Shows a chart of the 20 managed disks that are most active, and details for all managed disks, for a time period that you specify. Use the report to analyze the performance of managed disks.
- [Running the Performance of One Managed Disk](#)
To see performance metrics for managed disks on a storage system, use the Performance of One Managed Disk.
- [Performance of One Managed Disk report](#)
Shows four charts, and a table of performance metrics for a managed disk over a time period. Use the report to analyze the performance of a managed disk.
- [Running the Compare Performance of Multiple Managed Disks report](#)
Use the Compare Performance of Multiple Managed Disks report to compare up to four performance metrics for multiple managed disks that are on storage systems. For example, you can compare data rates, I/O rates, and response times.
- [Compare Performance of Multiple Managed Disks report](#)
Shows up to four performance metrics for multiple managed disks over time. Use the report to compare the performance metrics for managed disks that are on storage systems.

Running the Managed Disks Capacity report

To see information about used space and available space on managed disks (MDisks) on storage virtualizers, run the Managed Disks Capacity report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Managed Disks.
4. Click Managed Disks Capacity.
5. Optional: Select the configuration of the storage systems.
6. Select one or more storage systems.
7. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Managed Disks Capacity report](#)

Managed Disks Capacity report

Shows a chart of the 20 managed disks with the most available space, and a table with space statistics for all managed disks. Use the report to review space availability in your managed disks.

Restriction: This report applies only to SAN Volume Controller, Storwize®, and Hitachi Universal Storage Platform V systems.

Charts

The bar chart shows the used space and the available space for managed disk on storage systems.

Report output

For each managed disk on a storage system, the following information is provided:

Storage System MDisk Name

The name that was assigned to the managed disk on a storage system when it was added to the system.

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage System MDisk Capacity (GiB)

The amount of storage space on the managed disk.

Storage System MDisk Used Space (GiB)

The amount of used space on the managed disk.

Storage System MDisk Available Space (GiB)

The amount of storage space that is available on a managed disk. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage MDisk RAID Level

The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

Storage System MDisk Type

The type of managed disk on a storage system. For example, the disk on a storage system can be a local managed disk.

Storage System MDisk Mode

The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

Array

Extents are provided from local disks.

Managed

Extents are provided from a back-end storage volume.

Unmanaged

The managed disk is not used in the system.

Storage System MDisk Strip Size (KB)

The RAID strip size on a managed disk on a storage system.

Storage System MDisk Is Balanced

Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is **Yes**, the LUNs are balanced.

Storage System MDisk Fast Write State

Shows whether the cache for a volume on a disk that is managed by a storage system is empty, contains data, or is corrupted.

Storage System MDisk Write Verify

Shows whether all write operations on a managed disk on a storage system are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is **Yes**, all write operations are verified by a follow-up read operation.

Related tasks

- [Running the Managed Disks Capacity report](#)

Running the Most Active Managed Disks report

To see which managed disks are most active on storage systems, use the Most Active Managed Disks.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Managed Disks.
4. Click Most Active Managed Disks.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Managed Disks report](#)

Most Active Managed Disks report

Shows a chart of the 20 managed disks that are most active, and details for all managed disks, for a time period that you specify. Use the report to analyze the performance of managed disks.

Restriction: This report applies only to SAN Volume Controller, Storwize®, and Hitachi Universal Storage Platform V systems.

Charts

The bar chart shows the most active managed disks that are on storage systems. By default, managed disks are sorted by the I/O rate.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active managed disks is as follows:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage MDisk Name

The name that was assigned to the managed disk when it was added to the system.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Storage pool	Performance of One Pool
Managed disk	Performance of One Managed Disk, Storage Resource Relationships Summary

Related tasks

- [Running the Most Active Managed Disks report](#)

Running the Performance of One Managed Disk

To see performance metrics for managed disks on a storage system, use the Performance of One Managed Disk.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Managed Disks.
4. Click Performance of One Managed Disk.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a managed disk.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
10. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One Managed Disk report](#)

Performance of One Managed Disk report

Shows four charts, and a table of performance metrics for a managed disk over a time period. Use the report to analyze the performance of a managed disk.

Restriction: This report applies only to SAN Volume Controller, Storwize®, and Hitachi Universal Storage Platform V systems.

Charts

The charts show the following performance statistics for a managed disk on a storage system:

- The total I/O rates and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates and the write data rates
- The read cache percentages, the write cache percentages, and the volume utilization percentages

Report output

The following information is provided in the report for the period that you specified:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage MDisk Name

The name that was assigned to the managed disk when it was added to the system.

Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)

The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Queue Time (ms/op)

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Most Active Managed Disks
Storage pool	Performance of One Pool

Related tasks

- [Running the Performance of One Managed Disk](#)

Running the Compare Performance of Multiple Managed Disks report

Use the Compare Performance of Multiple Managed Disks report to compare up to four performance metrics for multiple managed disks that are on storage systems. For example, you can compare data rates, I/O rates, and response times.

Procedure

1. In the Welcome portal, click Team Content, IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Managed Disks.
4. Click Compare Performance of Multiple Managed Disks.
5. Select the storage systems.
6. Select managed disks. To select managed disks, complete the following steps:
 - a. Type a keyword.
For example, type `mdisk1` to find the names of managed disks that start with `mdisk1`.
To change the default search option, click Options.
Use the percent sign as a wildcard.
For example, type `a%` to find managed disks that begin with 'a' or 'A'. Type `%a` to find managed disks that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the managed disks on a resource.
 - b. Click Search.
If the search is successful, the managed disks are displayed in the Results list.
 - c. Select the managed disks that you require from the list, and then click Insert.
7. Click Next.
8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
9. Optional: To add more metrics to the report, complete the following steps:
 - a. Click Add Metric.
 - b. Select the category of performance metric, and then select the performance metric.
 - c. If you want to remove the last metric that you added, click Remove Metric.
10. Optional: Select an interval.
11. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
12. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple Managed Disks report](#)

Compare Performance of Multiple Managed Disks report

Shows up to four performance metrics for multiple managed disks over time. Use the report to compare the performance metrics for managed disks that are on storage systems.

Restriction: This report applies only to SAN Volume Controller, Storwize®, and Hitachi Universal Storage Platform V systems.

Charts

The charts show the performance metrics for the managed disks that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metric

Shows the time at which the performance metric data was collected.

Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

Managed disks

Shows the names of the managed disks, and the performance metric data for each managed disk.

Related tasks

- [Running the Compare Performance of Multiple Managed Disks report](#)

Predefined reports about storage pools

You can run predefined reports to analyze and compare the capacity and performance of storage pools.

- [Running the Pools Capacity report](#)
To see information about the used and available space of pools on storage systems, run the Pools Capacity report.
- [Pools Capacity report](#)
Shows a chart of the 20 pools with the most allocated space, and a table with space statistics for all pools. Use the report to review space allocation in your pools.
- [Running the Pools Historical Capacity report](#)
To see how space in pools on storage systems is used, run the Pools Historical Capacity report.
- [Pools Historical Capacity report](#)
Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.
- [Running the Most Active Pools report](#)
To see which pools are most active on storage systems, run the Most Active Pools report.
- [Most Active Pools report](#)
Shows a chart of the 20 storage pools that are most active, and details for all pools, for a time period that you specify. Use the report to analyze the performance of pools.
- [Running the Performance of One Pool report](#)
To see information about the performance of a pool on a storage system, run the Performance of One Pool report.
- [Performance of One Pool report](#)
Shows four charts, and a table of performance metrics for a storage pool over a time period that you specify. Use the report to analyze the performance of a storage pool.
- [Running the Compare Performance of Multiple Pools](#)
Use the Compare Performance of Multiple Pools report to compare up to four performance metrics for multiple pools that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.
- [Compare Performance of Multiple Pools report](#)
Shows up to four performance metrics for multiple storage pools over time. Use the report to compare the performance metrics for storage pools that are on storage systems.
- [Running the Compare Performance of One Pool over Time Ranges report](#)
To compare a performance metric for one pool in two date ranges, run the Compare Performance of One Pool over Time Ranges report.
- [Compare Performance of One Pool over Time Ranges report](#)
Shows one performance metric on one pool over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Running the Pools Capacity report

To see information about the used and available space of pools on storage systems, run the Pools Capacity report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Pools.
4. Click Pools Capacity.
5. Optional: Select the configuration of the storage systems.
6. Optional: Select one or more storage systems.
7. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Pools Capacity report](#)

Pools Capacity report

Shows a chart of the 20 pools with the most allocated space, and a table with space statistics for all pools. Use the report to review space allocation in your pools.

Charts

The bar chart shows the used space and available space for pools on storage systems.

Report output

For each pool on a storage system, the following information is provided:

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage Pool Capacity (GiB)

The total amount of storage space in a pool.

Storage Pool Used Space (GiB)

The amount of storage space that is used in the pool.

Storage Pool Available Space (GiB)

The amount of unused space that is not reserved for volumes in pools that are on the storage system. IBM Spectrum Control uses the following formula to determine this value:

$$\text{pool capacity} - \text{used space}$$

For XIV® systems and IBM Spectrum Accelerate, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Storage Pool Physical Allocation Percentage

The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

IBM Spectrum Control uses the following formula to determine the allocation percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Pool Virtual Allocation Percentage

The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

IBM Spectrum Control uses the following formula to determine the allocation percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes. An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Storage Pool Shortfall Percentage to estimate how critical the shortage of space is for a pool.

This value is only available for pools with thin-provisioned volumes.

Storage Pool Shortfall Percentage

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated.

The higher the percentage, the more critical the shortfall of pool space.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{virtual capacity} - \text{allocated space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage Pool Assigned Volume Space (GiB)

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Pool Unassigned Volume Space (GiB)

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Pool Total Volume Capacity (GiB)

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Pool Allocated Space (GiB)

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following resources:

- Resources other than those that are running IBM Spectrum Virtualize and are configured as back-end storage
- Resources that are running IBM Spectrum Virtualize, are configured as back-end storage, and are not thin-provisioned

Storage Pool Used Volume Space (GiB)

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For resources that are running IBM Spectrum Virtualize, you can preallocate thin-provisioned volume space when volumes are created. For these resources, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that IBM Spectrum Control collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Pool Real Configured Space (GiB)

The amount of storage space that is in an XIV or IBM Spectrum Accelerate pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage Pool Real Available Space (GiB)

The amount of unused space that is in an XIV or IBM Spectrum Accelerate pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage Pool RAID Level

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Pool Is Thin Provisioned

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is **Yes**, the resource is thin-provisioned.

Storage Pool Is Solid State

Shows whether there are solid-state drives in the pool. This property can contain the following values:

Mixed

The pool contains both hard disk drives and solid-state drives.

Non solid state

The pool contains no solid-state drives.

Solid state

The pool contains a least one solid-state drive.

Storage Pool Number of MDisks

The number of managed disks in a storage pool.

Storage Pool Extent Size (MiB)

The size of the extent that was specified when a pool was created. Smaller extent sizes limit the maximum size of the volumes that can be created in a pool. Smaller extent sizes minimize the amount of potentially wasted space per volume.

This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Pool Warning Level

The percentage of used capacity of the storage pool at which a warning is generated.

Storage Pool Repository Capacity (GiB)

The amount of space on all extents in the repository of a pool. This space can be used to allocate track space-efficient volumes. This attribute applies only to DS8000® storage systems.

Storage Pool Used Repository Space (GiB)

The amount of space on all extents in the repository of a pool that are allocated. This space can be used to allocate track space-efficient volumes. This attribute applies only to the DS8000 storage systems.

Report links

To open related reports, click the name of the storage pool and then click the name of the report. You can open the following related reports:

- Pools Historical Capacity
- Volumes Capacity

Related tasks

- [Running the Pools Capacity report](#)

Running the Pools Historical Capacity report

To see how space in pools on storage systems is used, run the Pools Historical Capacity report.

Before you begin

You can see how space in a pool on a storage system, or in pools in a storage resource group, are used.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Pools.
4. Click Pools Historical Capacity.
5. Select the type of historical space report that you want to create. You can select either of the following options:
 - Of a Pool in a Storage System
 - Click Next.
 - Optional: Select the configuration of the storage system.
 - Select a storage system, and then select a pool.
 - Of Pools in a Storage Resource Group
 - Click Next.

- Optional: Select the configuration of the storage system.
 - Select a storage resource group.
6. Select an interval.
 7. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
 8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Pools Historical Capacity report](#)

Pools Historical Capacity report

Shows the usage of space on storage pools over a specified period. Use the report to analyze changes in your storage pool requirements.

Charts

The chart shows the following space statistics for the pool on a storage system:

- The capacity of the pool
- The used space in the pool
- The available space in the pool
- The virtual volume capacity of the pool
- The used volume space in the pool
- The allocated space in the pool

The chart shows the following space statistics for the pool in the storage resource group:

- The capacity of the pools
- The used space in the pools
- The available space in the pools
- The virtual volume capacity of the pools
- The used volume space in the pools
- The allocated space in the pool

Report output

The following information is included in reports for a pool on a storage system or for pools in a storage resource group:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Date

The date that the data was collected for the report.

Storage System Pool Capacity (GiB)

The amount of storage space in pools that are on the storage system. For an XIV® or IBM Spectrum Accelerate, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Pool Available Space (GiB)

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

IBM Spectrum Control uses the following formula to determine this value:

pool capacity - used space

For XIV systems and IBM Spectrum Accelerate, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Storage Pool Allocated Space (GiB)

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following resources:

- Resources other than those that are running IBM Spectrum Virtualize and are configured as back-end storage
- Resources that are running IBM Spectrum Virtualize, are configured as back-end storage, and are not thin-provisioned

Storage Pool Used Space (GiB)

The amount of storage space that is used in the pool.

Storage Pool Total Volume Capacity (GiB)

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Pool Used Volume Space (GiB)

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For resources that are running IBM Spectrum Virtualize, you can preallocate thin-provisioned volume space when volumes are created. For these resources, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that IBM Spectrum Control collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Pool Assigned Volume Space (GiB)

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Pool Unassigned Volume Space (GiB)

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Pool Real Configured Space (GiB)

The amount of storage space that is in pools that have an associated soft size on an XIV or IBM Spectrum Accelerate. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System MDisk Available Space (GiB)

The amount of storage space that is available on a managed disk. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Pool Number of MDisk

The number of managed disks in a storage pool.

Storage Pool Number of Storage Volumes

The number of volumes in the storage pool.

Related tasks

- [Running the Pools Historical Capacity report](#)

Running the Most Active Pools report

To see which pools are most active on storage systems, run the Most Active Pools report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Pools.
4. Click Most Active Pools.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Pools report](#)

Most Active Pools report

Shows a chart of the 20 storage pools that are most active, and details for all pools, for a time period that you specify. Use the report to analyze the performance of pools.

Charts

The bar chart shows the most active pools on storage systems. By default, pools are sorted by the I/O rate.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active pools is as follows:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Storage pool	Performance of One Pool, Most Active Volumes, Most Active RAID Arrays, Storage Resource Relationships Summary

Related tasks

- [Running the Most Active Pools report](#)

Running the Performance of One Pool report

To see information about the performance of a pool on a storage system, run the Performance of One Pool report.

Procedure

1. In the Welcome portal, click Team Content, IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Pools.
4. Click Performance of One Pool.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.
7. Select a pool.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
10. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One Pool report](#)

Performance of One Pool report

Shows four charts, and a table of performance metrics for a storage pool over a time period that you specify. Use the report to analyze the performance of a storage pool.

Charts

The charts show the following performance metrics for a pool on a storage system:

- The total overall rates for I/O operations, and the total response times
- The overall I/O rates for read and write operations, and the response times for read and write operations
- The total I/O rates for back-end operations, and overall response times for back-end operations

- The I/O rates for back-end read operations and back-end write operations, and the response rates for back-end read and back-end write operations

Report output

The following information is provided in the report for the period that you specified:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Related reports

To open related reports, click the name of the storage system and then click the name of the report. You can open the following related reports:

- Performance of One Storage System
- Most Active Pools

Related tasks

- [Running the Performance of One Pool report](#)

Running the Compare Performance of Multiple Pools

Use the Compare Performance of Multiple Pools report to compare up to four performance metrics for multiple pools that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Pools.
4. Click Compare Performance of Multiple Pools.
5. Select the storage systems, and then click Refresh.
6. Select the pools, and then click Next.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
 - a. Click Add Metric.
 - b. Select the category of performance metric, and then select the performance metric.
 - c. If you want to remove the last metric that you added, click Remove Metric.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period.

It is easier to analyze small amounts of information over shorter periods.

Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
11. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple Pools report](#)

Compare Performance of Multiple Pools report

Shows up to four performance metrics for multiple storage pools over time. Use the report to compare the performance metrics for storage pools that are on storage systems.

Charts

The charts show the performance metrics for the storage pools that you selected for the period that you specified.

Report output

The following information is provided in the report:

Day, hour, or time

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Performance metrics

Shows the performance metric data for the resource.

Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

Pools

Shows the names of the storage pools, and the performance metric data for each storage pool.

Related tasks

- [Running the Compare Performance of Multiple Pools](#)

Running the Compare Performance of One Pool over Time Ranges report

To compare a performance metric for one pool in two date ranges, run the Compare Performance of One Pool over Time Ranges report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Pools.
4. Click Compare Performance of One Pool over Time Ranges.
5. Optional: Select the configuration of the storage systems.
6. Select the storage system.
7. Select the category of performance metric about which you want a report, and then select the performance metric.
8. Select an interval.
9. Specify the two date ranges or reporting periods that you want to compare.
For example, you can compare the performance metric for this week with last week.
Tip: To specify a start date and an end date for a reporting period, select Custom Date Range from the list of reporting periods.
10. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Compare Performance of One Pool over Time Ranges report](#)

Compare Performance of One Pool over Time Ranges report

Shows one performance metric on one pool over two time periods. For example, you can use this report to compare the I/O rate for this week and last week.

Charts

The chart shows the performance data of the pool on the performance metric you selected, over the time periods that you selected.

Report output

The following information is provided in the report:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Day in Range

The number of the day in the date range at which the performance metric data was collected, relative to the date range that you specified. The number of days that are shown is determined by the reporting period that was selected for the report. If Last 30 Days was selected, then entries for days 1 - 30 are shown on the report.

Hour in Range

The number of the hour in the day at which the performance metric data was collected. For example, 0 is between 12:00:00 a.m. and 12:59:59 a.m., 1 is between 01:00:00 a.m. and 01:59:59 a.m.

Minute in Range

The interval of minutes during which the performance metric data was collected. For example, 0 shows that performance data was collected 0 - 5 minutes after the hour. Similarly, 5 shows that the performance data that was collected 5 - 10 minutes after the hour.

First Date Range

The details of the performance metric data in the first date range. The details are displayed in two columns in the First Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the first date range.

Second Date Range

The details of the performance metric data in the second date range. The details are displayed in two columns in the Second Date Range column. Depending on the interval that you selected, the Time, Day, or Hour column shows the time at which the performance metric data was collected. The column with the performance metric name shows the data for the performance metric for the second date range.

Related tasks

- [Running the Compare Performance of One Pool over Time Ranges report](#)

Predefined reports about ports

You can run predefined reports to analyze and compare the performance of ports.

- [Running the Most Active Ports report](#)
To see which ports are most active on storage systems, run the Most Active Ports report.
- [Most Active Ports report](#)
Shows a chart of the 20 ports that are most active, and details for all ports, for a time period that you specify. Use the report to analyze the performance of ports.
- [Running the Performance of One Port report](#)
To see performance metrics for a port on a storage system, run the Performance of One Port report.
- [Performance of One Port report](#)
Shows four charts, and a table of performance metrics for a port over a time period that you specify. Use the report to analyze the performance of a port.
- [Running the Compare Performance of Multiple Ports report](#)
Use the Compare Performance of Multiple Ports report to compare up to four performance metrics for multiple ports that are on storage systems. For example, you can compare data rates, I/O rates, and bandwidth percentages.
- [Compare Performance of Multiple Ports report](#)
Shows up to four performance metrics for multiple ports over time. Use the report to compare the performance metrics for ports that are on storage systems.

Running the Most Active Ports report

To see which ports are most active on storage systems, run the Most Active Ports report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content, IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Ports.
4. Click Most Active Ports.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.

8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Ports report](#)

Most Active Ports report

Shows a chart of the 20 ports that are most active, and details for all ports, for a time period that you specify. Use the report to analyze the performance of ports.

Charts

The bar chart shows the most active ports on storage systems. By default, ports are sorted by the I/O rate.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active ports is as follows:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Port Name

The name that was assigned to the storage port when the storage system was added to the system.

Storage System Port Location

The physical location of a port on DS8000® storage systems and XIV® systems. For all other systems, this property shows the WWPN of the port.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Overall Port Response Time (ms/op)

The average number of milliseconds to complete a send or receive operation.

Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Overall Port Bandwidth Percentage

The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Number of the Fibre Channel (FC) port	Performance of One Port

Related tasks

- [Running the Most Active Ports report](#)

Running the Performance of One Port report

To see performance metrics for a port on a storage system, run the Performance of One Port report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Ports.
4. Click Performance of One Port.
5. Optional: Select the configuration of the storage systems.
6. Select a storage system.

7. Select a port.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
10. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One Port report](#)

Performance of One Port report

Shows four charts, and a table of performance metrics for a port over a time period that you specify. Use the report to analyze the performance of a port.

Charts

The charts show the following performance metrics for a port on a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, and the response times
- The read data rates, the write data rates, and the total data rate
- The send bandwidth percentage, the receive bandwidth percentage, and the overall bandwidth percentage

Report output

The following information is provided in the report for the period that you specified:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage Port Name

The name that was assigned to the storage port when the storage system was added to the system.

Storage System Port Location

The physical location of a port on DS8000® storage systems and XIV® systems. For all other systems, this property shows the WWPN of the port.

Port Send I/O Rate (ops/s)

The average number of I/O operations per second for operations in which data is sent from a port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive I/O Rate (ops/s)

The average number of I/O operations per second for operations in which the port receives data. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

Port Send Response Time (ms/op)

The average number of milliseconds to complete a send operation.

Port Receive Response Time (ms/op)

The average number of milliseconds to complete a receive operation.

Overall Port Response Time (ms/op)

The average number of milliseconds to complete a send or receive operation.

Port Send Data Rate (MiB/s)

The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)

The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Bandwidth Percentage

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage

The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Related reports

To open the Most Active Ports report, click the name of the storage system in the report table.

Related tasks

- [Running the Performance of One Port report](#)

Running the Compare Performance of Multiple Ports report

Use the Compare Performance of Multiple Ports report to compare up to four performance metrics for multiple ports that are on storage systems. For example, you can compare data rates, I/O rates, and bandwidth percentages.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Ports.
4. Click Compare Performance of Multiple Ports.
5. Select the storage systems.
6. Select ports. To select ports, complete the following steps:
 - a. Type a keyword.
For example, type `port1` to find the names of ports that start with `port1`.
To change the default search option, click Options.
Use the percent sign as a wildcard.
For example, type `a%` to find ports that begin with 'a' or 'A'. Type `%a` to find ports that begin with or contain the letter 'a' or 'A'. You can type `*` to retrieve all of the ports on a resource.
 - b. Click Search.
If the search is successful, the ports are displayed in the Results list.
 - c. Select the ports that you require from the list, and then click Insert.
7. Click Next.
8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
11. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple Ports report](#)

Compare Performance of Multiple Ports report

Shows up to four performance metrics for multiple ports over time. Use the report to compare the performance metrics for ports that are on storage systems.

Charts

The charts show the performance metrics for the ports that you selected for the period that you specified.

Report output

The following information is provided in the report:

Day, hour, or time

The day, hour, or time at which the information about the specified resource or device was collected. Whether the day, hour, or time is displayed depends on the interval that you select for the report.

Performance metrics

Shows the performance metric data for the resource.

Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

Ports

Shows the names of the storage ports, and the performance metric data for each port.

Related tasks

- [Running the Compare Performance of Multiple Ports report](#)

Predefined reports about RAID arrays

You can run predefined reports to analyze and compare the performance of RAID arrays.

- [Running the Most Active RAID Arrays report](#)
To see which arrays are most active on DS8000® storage systems, run the Most Active RAID Arrays report.
- [Most Active RAID Arrays report](#)
Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.
- [Running the Performance of One RAID Array report](#)
To see performance metrics for a RAID array on DS8000 storage systems, run the Performance of One RAID Array report.
- [Performance of One RAID Array report](#)
Shows four charts, and a table of performance metrics for a RAID array over a time period. Use the report to analyze the performance of a RAID array.
- [Running the Compare Performance of Multiple RAID Arrays report](#)
To compare the performance of multiple arrays that are on storage systems, use the Compare Performance of Multiple RAID Arrays report. For example, you can compare data rates, I/O rates, and response times.
- [Compare Performance of Multiple RAID Arrays report](#)
Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

Running the Most Active RAID Arrays report

To see which arrays are most active on DS8000® storage systems, run the Most Active RAID Arrays report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click RAID Arrays.
4. Click Most Active RAID Arrays.
5. Select the storage systems, and then select a sort order.
6. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
7. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active RAID Arrays report](#)

Most Active RAID Arrays report

Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

Restriction: This report applies only to DS8000® systems.

Charts

The bar chart shows the most active RAID on DS8000, DS6000™, and TotalStorage™ Enterprise Storage Server® storage systems. By default, arrays are sorted by the I/O rate.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active RAID is as follows:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Controller Name

The name that was assigned to the controller when it was added to the system.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage RAID Array Name

The name that was assigned to the RAID array when it was added to the system.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Disk Utilization Percentage

The average percentage of time that the disks that are associated with an array are busy.

No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value **N/A** is displayed. This limitation applies only to DS8000 storage systems.

If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.

Tip: Some highly sequential workloads such as batch or backup processing might continually exceed the threshold because they drive the arrays to high utilization percentages. For these types of workloads, a high utilization indicates that the work is being performed very efficiently and is not a cause for concern.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Controller	Performance of One Controller or Module
RAID array	Performance of One RAID Array

Related tasks

- [Running the Most Active RAID Arrays report](#)

Running the Performance of One RAID Array report

To see performance metrics for a RAID array on DS8000® storage systems, run the Performance of One RAID Array report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click RAID Arrays.
4. Click Performance of One RAID Array.
5. Select a storage system.
6. Select a RAID array.
7. Optional: Select an interval.
8. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
9. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One RAID Array report](#)

Performance of One RAID Array report

Shows four charts, and a table of performance metrics for a RAID array over a time period. Use the report to analyze the performance of a RAID array.

Restriction: This report applies only to DS8000® systems.

Charts

The charts show the following performance statistics for arrays on storage systems:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The back-end I/O rates, and the back-end I/O response times
- The back-end read I/O rates, the back-end write I/O rates, the read response times, and the write response times

Report output

The following information is provided in the report for the period that you specified:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Controller Name

The name that was assigned to the controller when it was added to the system.

Storage RAID Array Name

The name that was assigned to the RAID array when it was added to the system.

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Disk Utilization Percentage

The average percentage of time that the disks that are associated with an array are busy.

No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value **N/A** is displayed. This limitation applies only to DS8000 storage systems.

If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.

Tip: Some highly sequential workloads such as batch or backup processing might continually exceed the threshold because they drive the arrays to high utilization percentages. For these types of workloads, a high utilization indicates that the work is being performed very efficiently and is not a cause for concern.

For DS8000 storage systems, this property is available only for 8.5.0 and later. For earlier versions, the value **N/A** is shown.

Related reports

To open the Most Active RAID Arrays report, click the name of the storage system in the report table.

Related tasks

- [Running the Performance of One RAID Array report](#)

Running the Compare Performance of Multiple RAID Arrays report

To compare the performance of multiple arrays that are on storage systems, use the Compare Performance of Multiple RAID Arrays report. For example, you can compare data rates, I/O rates, and response times.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click RAID Arrays.
4. Click Compare Performance of Multiple RAID Arrays.
5. Select the storage systems, and then click Refresh.
6. Select the RAID, and then click Next.
7. Select the category of performance metric that you want to include in the report, and then select the performance metric.
8. Optional: To add more metrics to the report, complete the following steps:
 - a. Click Add Metric.
 - b. Select the category of performance metric, and then select the performance metric.
 - c. If you want to remove the last metric that you added, click Remove Metric.
9. Optional: Select an interval.
10. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
11. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple RAID Arrays report](#)

Compare Performance of Multiple RAID Arrays report

Shows up to four performance metrics for multiple arrays over time. Use the report to compare the performance metrics for arrays that are on storage systems.

Restriction: This report applies only to DS8000® systems.

Charts

The charts show the performance metrics for the arrays that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metric

Shows the time at which the performance metric data was collected.

Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

RAID arrays

Shows the names of the RAID arrays, and the performance metric data for each RAID array.

Related tasks

- [Running the Compare Performance of Multiple RAID Arrays report](#)

Predefined reports about volumes

You can run predefined reports to analyze the capacity and performance of volumes.

- [Running the Volumes Capacity report](#)
To see information about the capacity of volumes on storage systems, run the Volumes Capacity report.
- [Volumes Capacity report](#)
Shows a chart of the 20 volumes with the most used space, and a table with details of all volumes. Use the report to review the space allocation in your volumes.
- [Running the Volumes Historical Capacity report](#)
To see how space on volumes on storage systems is used, run the Volumes Historical Capacity report.
- [Volumes Historical Capacity report](#)
Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.
- [Running the Most Active Volumes report](#)
To see which volumes are most active on storage systems, run the Most Active Volumes report.
- [Most Active Volumes report](#)
Shows a chart of the 20 volumes that are most active, and details for all volumes, for a time period that you specify. Use the report to analyze the performance of

volumes.

- [Running the Performance of One Volume report](#)

To see performance metrics for volumes on a storage system, run the Performance of One Volume report.

- [Performance of One Volume report](#)

Shows four charts, and a table of performance metrics for a volume over a time period that you specify. Use the report to analyze the performance of a volume.

- [Running the Compare Performance of Multiple Volumes report](#)

Use the Compare Performance of Multiple Volumes report to compare up to four performance metrics for multiple volumes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

- [Compare Performance of Multiple Volumes report](#)

Shows up to four performance metrics for multiple volumes over time. Use the report to compare the performance metrics for volumes that are on storage systems.

Running the Volumes Capacity report

To see information about the capacity of volumes on storage systems, run the Volumes Capacity report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.

2. Click Storage Systems.

3. Click Volumes.

4. Click Volumes Capacity.

5. Optional: Select the configuration of the storage systems.

6. Select one or more storage systems.

7. Optional: If you selected SAN Volume Controller or Storwize® V7000 storage systems, you can select the type of volume.

Option	Description
All Volumes	Select All, the default, to view information about all of the volumes on the selected storage systems.
Only Primary Volumes	Select Primary to view information about primary volumes on the selected storage systems.
Only Secondary Volumes	Select Secondary to view information about secondary volumes on the selected storage systems.

8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Volumes Capacity report](#)

Volumes Capacity report

Shows a chart of the 20 volumes with the most used space, and a table with details of all volumes. Use the report to review the space allocation in your volumes.

Charts

The bar chart shows the following statistics for volumes on storage systems:

- Used space.
- Unallocated space.
- Unused allocated space, that is, the value of Storage Volume Allocated Space (GiB) minus the value of Storage Volume Used Space (GiB).

Report output

For each volume on a storage system, the following information is provided:

Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage Volume WWN

The worldwide name of the volume.

Storage Volume Capacity (GiB)

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV® systems and IBM Spectrum Accelerate, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Volume Allocated Space (GiB)

The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following resources:

- Resources other than those that are running IBM Spectrum Virtualize and are configured as back-end storage
- Resources that are running IBM Spectrum Virtualize, are configured as back-end storage, and are not thin-provisioned

Storage Volume Used Space (GiB)

The amount of allocated space that is used by a volume.

For resources that are running IBM Spectrum Virtualize, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that IBM Spectrum Control collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Volume Unallocated Space (GiB)

The amount of space in a pool that is not reserved for a volume.

IBM Spectrum Control uses the following formula to determine this value:

capacity - allocated space

Available only for thin provisioned volumes.

Storage Volume Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Volume Physical Allocation Percentage

The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

IBM Spectrum Control uses the following formula to determine this value:

(allocated space ÷ capacity) × 100

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Volume Shortfall Percentage

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume.

The higher the percentage, the more critical the shortfall of space.

IBM Spectrum Control uses the following formula to determine this value:

(unallocatable volume space ÷ volume unallocated space) × 100

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This Storage Volume Shortfall Percentage represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

**100 GiB unallocated volume space - 70 GiB remaining
pool space ÷ 100 GiB unallocated volume space × 100**

Storage Volume Used Allocated Space Percentage

The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

IBM Spectrum Control uses the following formula to determine this value:

(volume used space ÷ volume allocated space) × 100

This property is available only for volumes on resources that are running IBM Spectrum Virtualize.

Storage Volume RAID Level

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Volume Is Assigned

Shows whether the volume is assigned to a server. If this value is **Yes**, the volume is assigned to a server.

Storage Volume Is Thin Provisioned

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is **Yes**, the resource is thin-provisioned.

Storage Volume Type

The type of storage volume. For example, the storage volume can be striped or sequential. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Grain Size (KiB)

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Warning Level

The percentage of volume capacity that is used at which a warning is generated. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Logical Subsystem

The logical subsystem (LSS) to which a volume or pool belongs.

Storage Volume Copy ID
The identifier of a volume copy.
Storage Volume Copy Is Primary
Shows whether the volume copy is the primary volume copy.

Related reports

To open the Volumes Historical Capacity report, click the name of the volume in the report table.

Related tasks

- [Running the Volumes Capacity report](#)

Running the Volumes Historical Capacity report

To see how space on volumes on storage systems is used, run the Volumes Historical Capacity report.

Before you begin

You can see how space on a volume on a storage system, or on volumes in a storage resource group, are used.

Procedure

1. In the Welcome portal, click Team Content, IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Volumes.
4. Click Volumes Historical Capacity.
5. Select the type of historical capacity report that you want to create. You can select either of the following options:
 - Of a Volume in a Storage System
 - Click Next.
 - Optional: Select the configuration of the storage system.
 - Select a storage system.
 - Type a keyword and click Search. For example, type lun, to find the names of volumes that start with lun. To change the default search option, click Options.
Tip: Use the percent sign (%) as a wildcard. For example, type a% to find volume names that begin with 'a' or 'A'. Type %a to find volume names that begin with or contain the letter 'a' or 'A'. You can type % to retrieve all of the volumes on a resource.
 - Select a volume.
 - Of Volumes in a Storage Resource Group
 - Click Next.
 - Select a storage resource group.
6. Select an interval.
7. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Volumes Historical Capacity report](#)

Volumes Historical Capacity report

Shows the usage of space on volumes over a specified period. Use the report to analyze changes in your volume requirements, particularly the increasing usage of space on thin-provisioned volumes.

Charts

The chart shows the following space statistics for the volume on a storage system:

- The capacity of the volume
- The used space on the volume
- The space on the volume that is used for system management
- The allocated space on the volume

The chart shows the following space statistics for the volumes in a storage resource group:

- The capacity of the volumes
- The used space on the volumes
- The space on the volumes that is used for system management
- The allocated space on the volume

If the volume is thin provisioned, the chart shows the total amount of virtual space that is allocated to the volume, and the total amount of real space that is used by the volume.

Report output

The following information is included in reports for a volume on a storage system, or for volumes in a storage resource group:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

Date

The date that the data was collected for the report.

Storage Volume Capacity (GiB)

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV® systems and IBM Spectrum Accelerate, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Volume Used Space (GiB)

The amount of allocated space that is used by a volume.

For resources that are running IBM Spectrum Virtualize, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that IBM Spectrum Control collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Volume Allocated Space (GiB)

The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following resources:

- Resources other than those that are running IBM Spectrum Virtualize and are configured as back-end storage
- Resources that are running IBM Spectrum Virtualize, are configured as back-end storage, and are not thin-provisioned

Storage Volume Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Related tasks

- [Running the Volumes Historical Capacity report](#)

Running the Most Active Volumes report

To see which volumes are most active on storage systems, run the Most Active Volumes report.

Before you begin

To run the report with the default settings, click Finish. The report shows information about all of the resources in the list.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Volumes.
4. Click Most Active Volumes.
5. Optional: Select the configuration of the storage systems.
6. Select the storage systems, and then select a sort order.
7. To set the time frame of the report, specify a reporting period.
If the report does not include a particular resource that you want to see, select a reporting period for which data is available for the resource.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
8. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Most Active Volumes report](#)

Most Active Volumes report

Shows a chart of the 20 volumes that are most active, and details for all volumes, for a time period that you specify. Use the report to analyze the performance of volumes.

Charts

The bar chart shows the most active volumes on storage systems. By default, volumes are sorted by the I/O rate.

Tip: To change the sort order, select a value from the Sort Order list.

Report output

The information that is provided about the most active volumes is as follows:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Volume Utilization

The average percentage of time that the volume is busy.

Related reports

To open other reports that are related to a resource in this report, click the name of the resource. You can open the following related reports for the resources in this report:

Resource	Related reports
Storage system	Performance of One Storage System
Storage pool	Performance of One Pool
Volume	Performance of One Volume, Storage Resource Relationships Summary

Related tasks

- [Running the Most Active Volumes report](#)

Running the Performance of One Volume report

To see performance metrics for volumes on a storage system, run the Performance of One Volume report.

Procedure

1. In the Welcome portal, click Team Content, IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Volumes.
4. Click Performance of One Volume.
5. Optional: Select the configuration of the storage systems.

6. Select a storage system.
7. Select a volume. To select a volume, complete the following steps:
 - a. Type a keyword.
For example, type `1un`, to find the names of volumes that start with `1un`.
To change the default search option, click Options.
Use the percent sign as a wildcard.
For example, type `a%` to find volume names that begin with 'a' or 'A'. Type `%a` to find volume names that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the volumes on a resource.
 - b. Click Search.
If the search is successful, the names of the volumes are displayed in the Choice list.
 - c. From the list, select a volume.
8. Optional: Select an interval.
9. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
10. Click Finish.

Results

If information is available for the resources that you selected, a report is created. If no information is available for the selected resources, or data collection was unsuccessful, a message is displayed in the report.

Related reference

- [Performance of One Volume report](#)

Performance of One Volume report

Shows four charts, and a table of performance metrics for a volume over a time period that you specify. Use the report to analyze the performance of a volume.

Charts

The charts show the following performance metrics for a volume on a storage system:

- The total I/O rates, and the total response times
- The read I/O rates, the write I/O rates, the read response times, and the write response times
- The read data rates, and the write data rates
- The read cache percentages, the write cache percentages, and the volume utilization percentages

Report output

The following information is provided in the report for the period that you specified:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Volume Utilization

The average percentage of time that the volume is busy.

Related reports

To open the Most Active Volumes report, click the name of the storage system in the report table.

Related tasks

- [Running the Performance of One Volume report](#)

Running the Compare Performance of Multiple Volumes report

Use the Compare Performance of Multiple Volumes report to compare up to four performance metrics for multiple volumes that are on storage systems. For example, you can compare cache percentages, data rates, and I/O rates.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Predefined Reports.
2. Click Storage Systems.
3. Click Volumes.
4. Click Compare Performance of Multiple Volumes.
5. Select the storage systems.
6. Select volumes. To select volumes, complete the following steps:
 - a. Type a keyword.
For example, type `example1` to find the names of volumes that start with `example1`.
To change the default search option, click Options.
Use the percent sign as a wildcard.
For example, type `a%` to find volumes that begin with 'a' or 'A'. Type `%a` to find volumes that begin with or contain the letter 'a' or 'A'. You can type `%` to retrieve all of the volumes on a resource.
 - b. Click Search.
If the search is successful, the volumes are displayed in the Results list.
 - c. Select the volumes that you require from the list, and then click Insert.
7. Click Next.
8. Select the category of performance metric that you want to include in the report, and then select the performance metric.
9. Optional: To add more metrics to the report, complete the following steps:
 - a. Click Add Metric.
 - b. Select the category of performance metric, and then select the performance metric.
 - c. If you want to remove the last metric that you added, click Remove Metric.
10. Optional: Select an interval.
11. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
12. Click Finish.

Results

If the resources that you selected are not related, or if no data is available, a message is displayed in the report.

If information is not available for any of the resources that you selected, those resources are not included in the report.

Related reference

- [Compare Performance of Multiple Volumes report](#)

Compare Performance of Multiple Volumes report

Shows up to four performance metrics for multiple volumes over time. Use the report to compare the performance metrics for volumes that are on storage systems.

Charts

The charts show the performance metrics for the volumes that you selected for the period that you specified.

Report output

The following information is provided in the report:

Performance metrics

Shows the data for each performance metric.

Storage system names

Shows the names of the storage systems, and the performance metric data for each storage system.

Storage volume names

Shows the names of the storage volumes, and the performance metric data for each volume.

Related tasks

- [Running the Compare Performance of Multiple Volumes report](#)

Predefined reports about storage tiering

Use storage tier reports to help you determine the best storage-tier configuration for your environment.

Information collection

IBM Spectrum® Control collects information from:

- SAN Volume Controller
- Storwize® V7000
- Storwize V7000 Unified

IBM Spectrum Control provides reports that show performance and usage data for the SAN Volume Controller, Storwize V7000, and Storwize V7000 Unified systems in your environment.

Workload activity reports

Before you run workload activity reports, you must set the back-end properties of the managed disk (MDisk) groups in your environment.

You can use the IBM Spectrum Control CLI called `tcptool` to set the back-end properties of MDisk groups. Alternatively, you can set the properties in the MDisk Group Details page in the IBM Spectrum Control GUI.

With the information that is collected from IBM Spectrum Control, you can run workload activity reports. The reports provide information about the workload activity of these resources:

- Managed disk (MDisk) groups
- Virtual disks (VDisks) in MDisk groups
- VDisks in storage resource groups
- VDisks in storage virtualizers

By analyzing workload activity reports, you can determine:

- Whether the workload activity of an MDisk group is too low or too high
- Whether a VDisk should be migrated in one of the following ways:
 - To an MDisk group with a faster read I/O capability rate
 - To an MDisk group with a slower read I/O capability rate

Explanation: In the predefined reports about performance, resources, and historical information, MDisk groups are called pools, and VDisks are called volumes.

- [Running the Storage Resource Group - VDisk Workload Activity Details report](#)
To see which virtual disks (VDisks) are most active in a storage resource group, run the Storage Resource Group - VDisk Workload Activity Details report.
- [Storage Resource Group - VDisk Workload Activity Details report](#)
Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.
- [Running the VDisk Details report](#)
To see information about a virtual disk (VDisk), run the VDisk Details report.
- [VDisk Details report](#)
Shows the workload activity of a virtual disk (VDisk) in a managed disk group. Use the report to analyze the performance of a VDisk in a managed disk (MDisk) group.
- [Running the MDisk Groups - VDisk Workload Activity report](#)
To see which virtual disks (VDisks) are most active in managed disk (MDisk) groups, run the MDisk Groups - VDisk Workload Activity report.
- [MDisk Groups - VDisk Workload Activity report](#)
Use this report to monitor the performance of virtual disks (VDisks) in managed disk (MDisk) groups. The workload activity of each MDisk group is determined by calculating the average and the maximum peak utilization of the VDisks in the MDisk group.
- [Running the Storage Virtualizer - VDisk Workload Activity report](#)
To see which virtual disks (VDisks) are most active on a storage virtualizer, run the Storage Virtualizer - VDisk Workload Activity report.
- [Storage Virtualizer - VDisk Workload Activity report](#)
Shows the workload activity of virtual disks (VDisks) in a storage virtualizer. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.
- [Running the Storage Resource Group - VDisk Workload Activity report](#)
To see which virtual disks (VDisks) are most active in storage resource groups, run the Storage Resource Group - VDisk Workload Activity report.
- [Storage Resource Group - VDisk Workload Activity report](#)
Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

- [Running the MDisk Group - VDisk Workload Activity Details report](#)
To see which of the virtual disks in a managed disk (MDisk) group are the most active, run the MDisk Group - VDisk Workload Activity Details report.
- [MDisk Group - VDisk Workload Activity Details report](#)
Shows the workload activity of virtual disks (VDisks) in a managed disk (MDisk) group. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.
- [Running the MDisk Group Details report](#)
To see statistics about the performance of a managed disk (MDisk) group, run the MDisk Group Details report.
- [MDisk Group Details report](#)
Shows the workload activity of a managed disk (MDisk) group. Use the report to analyze storage services for an MDisk group.
- [Running the MDisk Groups - Workload Activity report](#)
To see which managed disk (MDisk) groups are most active, run the MDisk Groups - Workload Activity report.
- [MDisk Groups - Workload Activity report](#)
Shows the workload activity of managed disk (MDisk) groups. You can use the information that the report provides to investigate which MDisk groups are being underused or overused.
- [Read I/O capability formula](#)
The read I/O capability formula is used by IBM Spectrum Control to estimate the workload capacity of managed disk (MDisk) groups.
- [Changing threshold values in storage tier reports](#)
Use the `setdscfg` command to change the threshold values in workload activity reports for managed disk groups and virtual disks.
- [Threshold values in storage tier reports](#)
Threshold values are used to assess the workload activity of managed disk (MDisk) groups and virtual disks (VDisks) in storage tier reports. To change a threshold value, you must know the value of the property key that corresponds with the threshold value that is used in the report.

Related reference

- [setarray](#)
- [autosetarray](#)
- [lsarray](#)

Running the Storage Resource Group - VDisk Workload Activity Details report

To see which virtual disks (VDisks) are most active in a storage resource group, run the Storage Resource Group - VDisk Workload Activity Details report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Packages.
2. Click Storage Tiering.
3. Click VDisk Workload Activity Reports.
4. Select Storage Resource Group - VDisk Workload Activity Details, then click Next.
5. Select a storage virtualizer, and then select a storage resource group.
6. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
7. Click Finish.

Related reference

- [Storage Resource Group - VDisk Workload Activity Details report](#)

Storage Resource Group - VDisk Workload Activity Details report

Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

In the IBM Spectrum® Control GUI, storage resource groups are available as general groups.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Reporting period

The name of the storage virtualizer, the name of the storage resource group, the date filter such as `Last 7 days`, and the reporting period are displayed in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is `myverylongmanageddiskgroupname123` in a chart, it is truncated and displayed as `myverylongmanageddiskgroupname...`

Two charts are displayed:

Most Active VDIs chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

average(vdisk_peak_utilization)

To calculate maximum VDisk utilization, the following formula is used:

maximum(vdisk_peak_utilization)

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:

$$86.30 + 50.01 + 79.41 = 215.72$$

2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

$$215.72 \div 3 = 71.91$$

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

Least Active VDIs chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active VDIs chart are also used in the least active VDIs chart.

The number of VDIs that a storage resource group uses determines how many VDIs are displayed in each chart, as shown in the table.

Number of VDIs that the storage resource group uses	Number of VDIs in the Most Active VDIs chart	Number of VDIs in the Least Active VDIs chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

Tip: Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of the MDisk group in the report to open the MDisk Group Details Report. To monitor the workload activity of VDIs for all storage resource groups in a storage virtualizer, run the Storage Resource Group - VDisk Workload Activity Report.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr:min.). The convention that is used to represent a threshold value is

a number followed by the percent sign and enclosed in brackets: {%. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: $(\text{maximum_response_time_value} \times \text{maximum_response_time_threshold_value})/100$. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

The following information is provided about VDIs in the storage resource group:

MDisk Group

The name of the MDisk group.

VDisk

The name of the VDisk.

VDisk Report Date

The date that the data was collected for the report.

VDisk Capacity (GB)

The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is \geq the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is $<$ the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)

The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)

The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)

The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)

The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)

The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of {%) Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)

The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)

The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)

The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of {%) Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)

The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.

VDisk Average Read Cache Hits (%)

The average read cache hits in percent.

VDisk Max Read Cache Hits (%)

The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of {%) Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of {%) Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of {%) Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)

The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks

- [Running the Storage Resource Group - VDisk Workload Activity Details report](#)

Running the VDisk Details report

To see information about a virtual disk (VDisk), run the VDisk Details report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Packages.
2. Click Storage Tiering.
3. Click VDisk Details.
4. Select a storage virtualizer, and then select a VDisk.
5. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
6. Click Finish.

Related reference

- [VDisk Details report](#)

VDisk Details report

Shows the workload activity of a virtual disk (VDisk) in a managed disk group. Use the report to analyze the performance of a VDisk in a managed disk (MDisk) group.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Reporting period

The name of the storage virtualizer, the name of the VDisk, the date filter such as `Last 7 days`, and the reporting period are displayed in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDisks are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: $(\text{maximum_response_time_value} \times \text{maximum_response_time_threshold_value}) / 100$. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

The following information is displayed about the VDisk:

MDisk Group

The name of the MDisk group.

VDisk

The name of the VDisk.

VDisk Report Date

The date that the data was collected for the report.

VDisk Capacity (GB)

The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is \geq the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is $<$ the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)

The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)

The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)

The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)

The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)

The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)

The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)

The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)

The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)

The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.

VDisk Average Read Cache Hits (%)

The average read cache hits in percent.

VDisk Max Read Cache Hits (%)

The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)

The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks

- [Running the VDisk Details report](#)

Running the MDisk Groups - VDisk Workload Activity report

To see which virtual disks (VDIs) are most active in managed disk (MDisk) groups, run the MDisk Groups - VDisk Workload Activity report.

Procedure

1. In the Welcome portal, click Team Content, IBM Spectrum Control Packages.
2. Click Storage Tiering.
3. Click VDisk Workload Activity Reports.
4. Select MDisk Groups - VDisk Workload Activity, then click Next.
5. Select a storage virtualizer.
6. To set the time frame of the report, specify a reporting period.

It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
7. Click Finish.

Related reference

- [MDisk Groups - VDisk Workload Activity report](#)

MDisk Groups - VDisk Workload Activity report

Use this report to monitor the performance of virtual disks (VDisks) in managed disk (MDisk) groups. The workload activity of each MDisk group is determined by calculating the average and the maximum peak utilization of the VDisks in the MDisk group.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Reporting period

The name of the storage virtualizer, the date filter such as `Last 7 days`, and the reporting period are displayed in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is `myverylongmanageddiskgroupname123` in a chart, it is truncated and displayed as `myverylongmanageddiskgroupname...`

Two charts are shown:

Most Active MDisk Groups chart

For each MDisk group two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

`average(vdisk_peak_utilization)`

To calculate maximum VDisk utilization, the following formula is used:

`maximum(vdisk_peak_utilization)`

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:

$$86.30 + 50.01 + 79.41 = 215.72$$

2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

$$215.72 \div 3 = 71.91$$

In the sample, the average peak VDisk utilization is 71.91.

For each MDisk group, the maximum VDisk utilization is also shown in the chart. The maximum VDisk utilization is the highest value for a VDisk in the `VDisk Utilization (%)` column. In the sample, the highest value for VDisk1 is 86.30.

Least Active MDisk Groups chart

For each MDisk group two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active MDisk group chart are also used in the least active MDisk group chart.

The number of MDisk groups in the storage virtualizer determines how many MDisk groups are shown in each chart, as shown in the following table.

Number of MDisk groups in the storage virtualizer	Number of MDisk groups in the Most Active MDisk Groups chart	Number of MDisk groups in the Least Active MDisk Groups chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

Tip: Click the bar that represents an MDisk group in a chart to open the MDisk Group - VDisk Workload Activity Details Report. Alternatively, click the name of the MDisk group in the report. Click the name of the VDisk in the report to open the VDisk Details Report.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: $(\text{maximum_response_time_value} \times \text{maximum_response_time_threshold_value}) / 100$. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

For each VDisk in an MDisk group, the following information is provided:

MDisk Group

The name of the MDisk group.

VDisk

The name of the VDisk.

VDisk Report Date

The date that the data was collected for the report.

VDisk Capacity (GB)

The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is ≥ the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is < the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)

The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)

The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)

The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)

The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)

The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)

The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)

The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)

The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of { } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)

The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.

VDisk Average Read Cache Hits (%)

The average read cache hits in percent.

VDisk Max Read Cache Hits (%)

The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of { } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of { } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of { } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)

The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks

- [Running the MDisk Groups - VDisk Workload Activity report](#)

Running the Storage Virtualizer - VDisk Workload Activity report

To see which virtual disks (VDIs) are most active on a storage virtualizer, run the Storage Virtualizer - VDisk Workload Activity report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Packages.
2. Click Storage Tiering.
3. Click VDisk Workload Activity Reports.
4. Select Storage Virtualizer - VDisk Workload Activity, then click Next.
5. Select a storage virtualizer.
6. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
7. Click Finish.

Related reference

- [Storage Virtualizer - VDisk Workload Activity report](#)

Storage Virtualizer - VDisk Workload Activity report

Shows the workload activity of virtual disks (VDIs) in a storage virtualizer. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Reporting period

The name of the storage virtualizer, the date filter such as Last 7 days, and the reporting period are shown in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is `myverylongmanageddiskgroupname123` in a chart, it is truncated and displayed as `myverylongmanageddiskgroupname...`

Two charts are shown:

Most Active VDIs chart

For each storage virtualizer, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

average(vdisk_peak_utilization)

To calculate maximum VDisk utilization, the following formula is used:

maximum(vdisk_peak_utilization)

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:

$$86.30 + 50.01 + 79.41 = 215.72$$

2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

$$215.72 \div 3 = 71.91$$

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

Least Active VDIs chart

For each storage virtualizer, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active VDIs chart are also used in the least active VDIs chart.

The number of VDIs in a storage virtualizer determines how many VDIs are shown in each chart, as shown in the following table.

Number of VDIs in the storage virtualizer	Number of VDIs in the Most Active VDIs chart	Number of VDIs in the Least Active VDIs chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

Tip: Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: $(\text{maximum_response_time_value} \times \text{maximum_response_time_threshold_value}) / 100$. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

For each VDisk in a storage virtualizer, the following information is shown:

MDisk Group

The name of the MDisk group.

VDisk

The name of the VDisk.

VDisk Report Date

The date that the data was collected for the report.

VDisk Capacity (GB)

The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is \geq the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is $<$ the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)

The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)

The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)

The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)

The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)

The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)

The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)

The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)

The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)

The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.

VDisk Average Read Cache Hits (%)

The average read cache hits in percent.

VDisk Max Read Cache Hits (%)

The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)
The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks

- [Running the Storage Virtualizer - VDisk Workload Activity report](#)

Running the Storage Resource Group - VDisk Workload Activity report

To see which virtual disks (VDisks) are most active in storage resource groups, run the Storage Resource Group - VDisk Workload Activity report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Packages.
2. Click Storage Tiering.
3. Click VDisk Workload Activity Reports.
4. Select Storage Resource Group - VDisk Workload Activity.
5. Select a storage virtualizer.
6. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
7. Click Finish.

Related reference

- [Storage Resource Group - VDisk Workload Activity report](#)

Storage Resource Group - VDisk Workload Activity report

Shows the workload activity of virtual disks (VDisks) in storage resource groups. The workload activity is determined by calculating the average and maximum peak utilization of the VDisks in the group.

In the IBM Spectrum® Control GUI, storage resource groups are available as general groups.
Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Reporting period

The name of the storage virtualizer, the date filter such as `Last 7 days`, and the reporting period are shown in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is `myverylongmanageddiskgroupname123` in a chart, it is truncated and displayed as `myverylongmanageddiskgroupname...`

Two charts are shown:

Most Active Storage Resource Groups chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

`average(vdisk_peak_utilization)`

To calculate maximum VDisk utilization, the following formula is used:

`maximum(vdisk_peak_utilization)`

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:

$$86.30 + 50.01 + 79.41 = 215.72$$

2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

$$215.72 \div 3 = 71.91$$

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

Least Active Storage Resource Groups chart

For each storage resource group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active storage resource groups chart are also used in the least active storage resource groups chart.

The number of storage resource groups in a storage virtualizer determines how many storage resource groups are shown in each chart, as shown in the table.

Number of storage resource groups in the storage virtualizer	Number of storage resource groups in the Most Active Storage Resource Groups chart	Number of storage resource groups in the Least Active Storage Resource Groups chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

Tip: Click the bar that represents a storage resource group to open the Storage Resource Group: VDisk Workload Activity Details Report. Alternatively, click the name of the storage resource group in the report. Click the name of the VDisk in the report to open the VDisk Details Report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: $(\text{maximum_response_time_value} \times \text{maximum_response_time_threshold value})/100$. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

For each storage resource group in a storage virtualizer, the following information is provided about VDIs:

Storage Resource Group

The name of the storage resource group.

MDisk Group

The name of the MDisk group.

VDisk

The name of the VDisk.

VDisk Report Date

The date that the data was collected for the report.

VDisk Capacity (GB)

The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is \geq the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is $<$ the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)

The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)

The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)

The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)

The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)

The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of { % } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)

The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)

The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)

The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of { % } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)

The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.

VDisk Average Read Cache Hits (%)

The average read cache hits in percent.

VDisk Max Read Cache Hits (%)

The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of { % } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of { % } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of { % } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)

The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks

- [Running the Storage Resource Group - VDisk Workload Activity report](#)

Running the MDisk Group - VDisk Workload Activity Details report

To see which of the virtual disks in a managed disk (MDisk) group are the most active, run the MDisk Group - VDisk Workload Activity Details report.

Procedure

1. In the Welcome portal, click Team Content \geq IBM Spectrum Control Packages.
2. Click Storage Tiering.
3. Click MDisk Group - VDisk Workload Activity Details.
4. Select a storage virtualizer, and then select an MDisk group.

5. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
6. Click Finish.

Related reference

- [MDisk Group - VDisk Workload Activity Details report](#)

MDisk Group - VDisk Workload Activity Details report

Shows the workload activity of virtual disks (VDisks) in a managed disk (MDisk) group. The workload activity is determined by calculating the average and maximum peak utilization of each VDisk.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Reporting period

The name of the storage virtualizer, the name of the MDisk group, the date filter such as `Last 7 days`, and the reporting period are displayed on the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

VDisk migration

If a VDisk is migrated to another MDisk group during a reporting period, daily VDisk reports are produced only for the target MDisk group. For the target MDisk group, VDisk reports are produced from the day that the VDisk is migrated from the source MDisk group until the last day of the reporting period.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is `myverylongmanageddiskgroupname123` in a chart, it is truncated and displayed as `myverylongmanageddiskgroupname...`

Two charts are shown:

Most Active VDisks chart

For each MDisk group, two bars are shown:

- A yellow bar that shows average VDisk utilization
- A red bar that shows maximum VDisk utilization

To calculate average VDisk utilization, the following formula is used:

average (vdisk_peak_utilization)

To calculate maximum VDisk utilization, the following formula is used:

maximum (vdisk_peak_utilization)

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

To calculate the average value of a VDisk:

1. The values for VDisk1 in the VDisk Utilization (%) column are added:

$$86.30 + 50.01 + 79.41 = 215.72$$

2. The total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

$$215.72 \div 3 = 71.91$$

In the sample, the average peak VDisk utilization is 71.91.

In this sample excerpt of report output for VDisk1, the values that are required to calculate average and maximum VDisk utilization are provided.

VDisk Report Date	VDisk Peak Utilization (%)
Dec 6 2011	86.30
Dec 7 2011	50.01
Dec 8 2011	79.41

Least Active VDisks chart

For each MDisk group, two bars are shown:

- A yellow bar that shows average VDisk utilization

- A red bar that shows maximum VDisk utilization

The same formulas that are used in the most active MDisk group chart are also used in the least active MDisk group chart.

The number of VDIs in an MDisk group determines how many VDIs are shown in each chart, as shown in the following table.

Number of VDIs in the MDisk group	Number of VDIs in the Most Active VDIs chart	Number of VDIs in the Least Active VDIs chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

Tip: Click the bar that represents a VDisk in a chart to open the VDisk Details Report. Alternatively, click the name of the VDisk in the report. Click the name of the MDisk group in the report to open the MDisk Group Details Report.

Threshold values

The names of the threshold values that are used in the report to monitor workload activity for VDIs are as follows:

- VDisk Max Read Response Time Threshold
- VDisk Max Write Response Time Threshold
- Max Read Cache Hits Time Threshold
- Max Read I/O Rate Threshold
- Max Write I/O Rate Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the VDisk maximum read response time threshold value is changed during a reporting period from 70% to 80%, the name of the column is changed to VDisk Max Read Response Time Threshold of {70 - 80%} Exceeded By (hr.:min.). The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used: $(\text{maximum_response_time_value} \times \text{maximum_response_time_threshold_value}) / 100$. The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

For each VDisk in an MDisk group, the following information is provided:

MDisk Group

The name of the MDisk group.

VDisk

The name of the VDisk.

VDisk Report Date

The date that the data was collected for the report.

VDisk Capacity (GB)

The total storage capacity of the VDisk in GB.

VDisk Thin Provisioning Capacity Status

The warning level threshold is set in the thin provisioning profile for the VDisk. The threshold is exceeded when the capacity of the VDisk that is used exceeds or is equal to the warning level threshold. If the used capacity of the VDisk in GB is ≥ the specified warning threshold value, a status value of `Warning` is displayed. If the used capacity of the VDisk in GB is < the specified warning threshold value, a status value of `OK` is displayed. If thin provisioning is not used, a status value of `Not Applicable` is displayed.

VDisk Peak Utilization (%)

The highest value of VDisk utilization over a 24-hour period.

VDisk Read I/O Rate (ops/s)

The read I/O rate of the VDisk in operations per second.

VDisk Max Read I/O Rate (ops/s)

The maximum read I/O rate of the VDisk in operations per second.

VDisk Average Read Response Time (ms/op)

The average read response time in milliseconds per operation.

VDisk Max Read Response Time (ms/op)

The maximum read response time in milliseconds per operation.

VDisk Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Write I/O Rate (ops/s)

The write I/O rate in operations per second.

VDisk Average Write Response Time (ms/op)

The average write response time in milliseconds per operation.

VDisk Max Write Response Time (ms/op)

The maximum write response time in milliseconds per operation.

VDisk Max Write Response Time Threshold of { } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write response time threshold for the VDisk is exceeded over a 24 hour period.

VDisk Max Read Density Rate (ops/s per GB)

The maximum read density rate in operations per second per GB. This value is calculated by dividing the maximum read I/O rate in operations per second by the number of GB that the VDisk was allocated.

VDisk I/O Rate Share of Total VDisk Read I/O Rate (%)

The maximum read I/O rate of the VDisk as a percentage of the total maximum read I/O rate of all the VDIs in the MDisk group.

VDisk Capacity Share of Total VDisk Capacity (%)

The real capacity of the VDisk as a percentage of the real capacity of all the VDIs in the MDisk group.

VDisk Average Read Cache Hits (%)

The average read cache hits in percent.

VDisk Max Read Cache Hits (%)

The maximum read cache hits in percent.

VDisk Max Read Cache Hits Time Threshold of { } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read cache hits time threshold is exceeded over a 24 hour period.

VDisk Max Read I/O Rate Threshold of { } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write I/O Rate Threshold of { } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write I/O rate threshold is exceeded over a 24 hour period.

VDisk Max Write Cache Delay I/O Rate (ops/s)

The maximum write cache delay I/O rate of the VDisk in operations per second.

Related tasks

- [Running the MDisk Group - VDisk Workload Activity Details report](#)

Running the MDisk Group Details report

To see statistics about the performance of a managed disk (MDisk) group, run the MDisk Group Details report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Packages.
2. Click Storage Tiering.
3. Click MDisk Group Details.
4. Select a storage virtualizer, and then select an MDisk group.
5. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
6. Click Finish.

Related reference

- [MDisk Group Details report](#)

MDisk Group Details report

Shows the workload activity of a managed disk (MDisk) group. Use the report to analyze storage services for an MDisk group.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Reporting period

The name of the storage virtualizer, the name of the MDisk group, the date filter such as Last 7 days, and the reporting period are displayed in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

Read I/O rate capability

If you do not set the back-end pool properties of an MDisk group:

- `Not Set` is displayed in the MDisk Group Read I/O Rate Capability column
- `Not Available` is displayed in the MDisk Group Read I/O Rate Availability column

If the read I/O rate capability of an MDisk group is less than the maximum read I/O rate of the MDisk group, negative values are displayed in the report.

Threshold values

The names of the threshold values in the report that are used to monitor workload activity for MDisk groups are:

- Actual Size Threshold
- Read I/O Rate Threshold
- Max Read Response Time Threshold
- Max Write Response Time Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the value of the actual size threshold is changed during a reporting period from 50% to 60%, the name of the column is changed to MDisk Group Available Size < Actual Size Threshold of {50 - 60%}. The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:

$$(\text{maximum_response_time_value} \times \text{maximum_response_time_threshold_value}) \div 100.$$

Tip: To monitor the workload activity of all MDisk groups in a storage virtualizer, run the MDisk Groups - Workload Activity report.

Report output

The following information is provided about the MDisk group:

MDisk Group

The name of the MDisk group.

MDisk Group Report Date

The date that the data was collected for the report.

MDisk Group Capacity Status

If the storage capacity that is allocated is less than the actual storage capacity, a status value of `OK` is displayed. If the storage capacity that is allocated is more than the actual storage capacity, a status value of `Warning` is displayed. The actual capacity of the MDisk group is calculated as follows:

$$(\text{mdisk_group_real_capacity_value} \div \text{mdisk_group_capacity_value}) \times 100$$

MDisk Group Capacity (GB)

The total storage capacity in GB of all of the disks in the MDisk group.

MDisk Group Real Capacity (GB)

The total amount in GB of real storage capacity in the MDisk group.

MDisk Group Free Space (GB)

The total amount of remaining capacity in GB that can be allocated to increase the size of an existing VDisk, or allocated to an additional VDisk.

MDisk Group Available Size < Actual Size Threshold of {%

If the amount of storage space that is available for allocation to VDIs is less than the actual size threshold, `Yes` is displayed in the column. If the amount of storage space that is available for allocation to VDIs is more than the actual size threshold, `No` is displayed in the column.

MDisk Group Read I/O Rate (ops/s)

The read I/O rate of the MDisk group in operations per second.

MDisk Group Read I/O Rate Capability (ops/s)

The read I/O capability rate of the MDisk group in operations per second.

Tip: To find out how the read I/O capability rate is calculated, see [Read I/O capability formula](#)

MDisk Group Max Read I/O Rate (ops/s)

The maximum read I/O rate of the MDisk group in operations per second.

MDisk Group Read I/O Rate Availability (ops/s)

The difference between the read I/O capability rate and the maximum read I/O rate in operations per second.

MDisk Group Read I/O Rate Availability < Read I/O Rate Capability Threshold of {%

The availability rate for read I/O operations is compared with the read I/O rate capability threshold. If the availability rate for read I/O operations is less than the read I/O rate capability threshold, `Yes` is displayed in the column. If the availability rate for read I/O operations is more than the read I/O rate capability threshold, `No` is displayed in the column.

MDisk Group Max Read Response Time (ms/op)

The maximum read response time in milliseconds per operation.

MDisk Group Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read response time threshold is exceeded over a 24 hour period.

MDisk Group Average Read Response Time (ms/op)

The average read response time in milliseconds per operation.

- MDisk Group Max Write I/O Rate (ops/s)
The maximum write I/O rate in operations per second.
- MDisk Group Max Write Response Time (ms/op)
The maximum write response time in milliseconds per operation.
- MDisk Group Max Write Response Time Threshold of { % } Exceeded By (hr.:min.)
The number of hours and minutes that the maximum write response time threshold is exceeded over a 24 hour period.
- MDisk Group Average Write Response Time (ms/op)
The average write response time in milliseconds per operation.
- MDisk Group Average Write I/O Size (KB)
The average size of write I/O operations in KB.

Related tasks

- [Running the MDisk Group Details report](#)

Running the MDisk Groups - Workload Activity report

To see which managed disk (MDisk) groups are most active, run the MDisk Groups - Workload Activity report.

Procedure

1. In the Welcome portal, click Team Content > IBM Spectrum Control Packages.
2. Click Storage Tiering.
3. Click MDisk Groups - Workload Activity.
4. Select a storage virtualizer.
5. To set the time frame of the report, specify a reporting period.
It is easier to analyze small amounts of information over shorter periods.
Tip: To specify a start date and an end date, select Custom date range from the list of reporting periods.
6. Click Finish.

Related reference

- [MDisk Groups - Workload Activity report](#)
- [Read I/O capability formula](#)

MDisk Groups - Workload Activity report

Shows the workload activity of managed disk (MDisk) groups. You can use the information that the report provides to investigate which MDisk groups are being underused or overused.

Restriction: This report applies only to SAN Volume Controller and Storwize® systems.

Reporting periods

The name of the storage virtualizer, the date filter such as `Last 7 days`, and the reporting period are displayed in the report.

Retention of report data

If you remove an MDisk group or a VDisk, the daily data that is collected for reports is retained for a configurable period. The default is 90 days. If you run a report after the data retention period expired, the daily data for the MDisk group or the VDisk is unavailable.

Read I/O rate capability

To estimate the read I/O rate capability of an MDisk group, you must set the back-end pool properties of MDisk groups.

If you do not set the back-end pool properties of an MDisk group:

- `Not Set` is displayed in the MDisk Group Read I/O Rate Capability column
- `Not Available` is displayed in the MDisk Group Read I/O Rate Availability column

If the read I/O rate capability of an MDisk group is less than the maximum read I/O rate of the MDisk group, negative values are displayed in the report.

Charts

The maximum length of an MDisk group, VDisk, or storage resource group name in charts is 25 characters. Names that are longer than 25 characters are truncated. To indicate that a name is truncated, an ellipsis (...) is appended to the name. For example, if the name of an MDisk group is `myverylongmanageddiskgroupname123` in a chart, it is truncated and displayed as `myverylongmanageddiskgroupname...`

Two charts are displayed:

- Most Active MDisk Groups
- Most Available MDisk Groups

Most Active MDisk Groups

For each MDisk group two bars are used:

- A yellow bar that shows the average activity value
- A red bar that shows the maximum activity value

To calculate the average activity value of an MDisk group, the following formula is used:

average(maximum_read_I/O_rate ÷ read_I/O_rate_capability)

To calculate the maximum activity value, the following formula is used:

maximum(maximum_read_I/O_rate ÷ read_I/O_rate_capability)

In this sample excerpt of report output for mdiskgroup1, the values that are required to calculate average and maximum activity values are provided.

MDisk Group Report Date	MDisk Group Read I/O Rate Capability (ops/s)	MDisk Group Max Read I/O Rate (ops/s)
Dec 6 2011	331	216.77
Dec 7 2011	331	256.40
Dec 8 2011	331	270.87

The average activity value of an MDisk group is calculated as follows:

1. For each report in the MDisk group, the value in the MDisk Group Max Read I/O Rate column is divided by the value in the MDisk Group Read I/O Rate Capability column:

$$\begin{aligned} 216.77 \div 331 &= 0.65 \\ 256.40 \div 331 &= 0.77 \\ 270.87 \div 331 &= 0.81 \end{aligned}$$

2. The result of the operation for each report is added:

$$0.65 + 0.77 + 0.81 = 2.23$$

3. To get the average value, the total is divided by the number of days that are specified in the date range. For example, if the reporting period spans 3 days, the total is divided by 3.

$$2.23 \div 3 = 0.74$$

The maximum value of an MDisk group is calculated as follows. For each report in an MDisk group, the value in the MDisk Group Max Read I/O Rate column is divided by the value in the MDisk Group Read I/O Rate Capability column:

$$\begin{aligned} 216.77 \div 331 &= 0.65 \\ 256.40 \div 331 &= 0.77 \\ 270.87 \div 331 &= 0.81 \end{aligned}$$

In this sample, the maximum value is 0.81.

Most Available MDisk Groups

For each MDisk group two bars are used:

- A blue bar to represent minimum values
- A red bar to represent average values

In this sample excerpt of report output for mdiskgroup2, the values that are required to calculate minimum and average availability values for MDisk Groups are provided.

MDisk Group Report Date	MDisk Group Read I/O Rate Availability (ops/s)
Dec 6 2011	435.41
Dec 7 2011	560.92
Dec 8 2011	480.95

The minimum value of an MDisk group is calculated as follows:

- For all of the reports for an MDisk group, the lowest value in the MDisk Group Read I/O Rate Availability column is used. In this sample, the lowest value for mdiskgroup2 is 435.41.

The average value of an MDisk group is calculated as follows:

1. For each report in an MDisk group, the values in the MDisk Group Read I/O Rate Availability column are added:

$$435.41 + 560.92 + 480.95 = 1477.28$$

2. To get the average value, the total is divided by the number of days that are specified in the date range. (For example, if the reporting period spans 3 days, the total is divided by 3.)

$$1477.28 \div 3 = 492.42$$

In this sample, the average value is 492.42.

The number of MDisk groups in a storage virtualizer determines how many MDisk groups are displayed in each chart, as shown in the following table.

Number of MDisk groups in the storage virtualizer	Number of MDisk groups in the Most Active MDisk Groups chart	Number of MDisk groups in the Least Active MDisk Groups chart
≥ 10	5	5
9	5	4
8	4	4
7	4	3
6	3	3

Number of MDisk groups in the storage virtualizer	Number of MDisk groups in the Most Active MDisk Groups chart	Number of MDisk groups in the Least Active MDisk Groups chart
5	3	2
4	2	2
3	2	1
2	1	1
1	1	0

Tip: Click the bar that represents an MDisk group in a chart to open the MDisk Group - VDisk Workload Activity Details Report. Alternatively, you can click the name of the MDisk group in the report.

Threshold values

The names of the threshold values in the report that are used to monitor workload activity for MDisk groups are:

- Actual Size Threshold
- Read I/O Rate Threshold
- Max Read Response Time Threshold
- Max Write Response Time Threshold

Threshold values are set during installation. If a threshold value is changed during a reporting period, the original value and the new threshold value are displayed as minimum and maximum values. If, for example, the value of the actual size threshold is changed during a reporting period from 50% to 60%, the name of the column is changed to MDisk Group Available Size < Actual Size Threshold of {50 - 60%}. The convention that is used to represent a threshold value is a number followed by the percent sign and enclosed in brackets: {%}. When the report is run, the variable value is replaced with the threshold value. To calculate whether read or write response time thresholds are exceeded, the following formula is used:

$$(\text{maximum_response_time_value} \times \text{maximum_response_time_threshold_value}) \div 100.$$

The result is converted from seconds to hours and minutes, and displayed in the report. Read and write response times are sampled at 5-minute intervals over a 24-hour period.

Report output

For each MDisk group, the following information is provided:

MDisk Group

The name of the MDisk group.

MDisk Group Report Date

The date that the data was collected for the report.

MDisk Group Capacity Status

If the storage capacity that is allocated is less than the actual storage capacity, a status value of **OK** is displayed. If the storage capacity that is allocated is more than the actual storage capacity, a status value of **Warning** is displayed. The actual capacity of the MDisk group is calculated as follows:

$$(\text{mdisk_group_real_capacity_value} \div \text{mdisk_group_capacity_value}) \times 100$$

MDisk Group Capacity (GB)

The total storage capacity in GB of all of the disks in the MDisk group.

MDisk Group Real Capacity (GB)

The total amount in GB of real storage capacity in the MDisk group.

MDisk Group Free Space (GB)

The total amount of remaining capacity in GB that can be allocated to increase the size of an existing VDisk, or allocated to an additional VDisk.

MDisk Group Available Size < Actual Size Threshold of {%}

If the amount of storage space that is available for allocation to VDIs is less than the actual size threshold, **Yes** is displayed in the column. If the amount of storage space that is available for allocation to VDIs is more than the actual size threshold, **No** is displayed in the column.

MDisk Group Read I/O Rate (ops/s)

The read I/O rate of the MDisk group in operations per second.

MDisk Group Read I/O Rate Capability (ops/s)

The read I/O capability rate of the MDisk group in operations per second.

Tip: To find out how the read I/O capability rate is calculated, see [Read I/O capability formula](#)

MDisk Group Max Read I/O Rate (ops/s)

The maximum read I/O rate of the MDisk group in operations per second.

MDisk Group Read I/O Rate Availability (ops/s)

The difference between the read I/O capability rate and the maximum read I/O rate in operations per second.

MDisk Group Read I/O Rate Availability < Read I/O Rate Capability Threshold of {%}

The availability rate for read I/O operations is compared with the read I/O rate capability threshold. If the availability rate for read I/O operations is less than the read I/O rate capability threshold, **Yes** is displayed in the column. If the availability rate for read I/O operations is more than the read I/O rate capability threshold, **No** is displayed in the column.

MDisk Group Max Read Response Time (ms/op)

The maximum read response time in milliseconds per operation.

MDisk Group Max Read Response Time Threshold of {%} Exceeded By (hr.:min.)

The number of hours and minutes that the maximum read response time threshold is exceeded over a 24 hour period.

MDisk Group Average Read Response Time (ms/op)

The average read response time in milliseconds per operation.

MDisk Group Max Write I/O Rate (ops/s)

The maximum write I/O rate in operations per second.

MDisk Group Max Write Response Time (ms/op)

The maximum write response time in milliseconds per operation.

MDisk Group Max Write Response Time Threshold of { } Exceeded By (hr.:min.)

The number of hours and minutes that the maximum write response time threshold is exceeded over a 24 hour period.

MDisk Group Average Write Response Time (ms/op)

The average write response time in milliseconds per operation.

MDisk Group Average Write I/O Size (KB)

The average size of write I/O operations in KB.

Related tasks

- [Running the MDisk Groups - Workload Activity report](#)

Read I/O capability formula

The read I/O capability formula is used by IBM Spectrum® Control to estimate the workload capacity of managed disk (MDisk) groups.

The estimated workload capacity is compared with the actual workload capacity to assess the most active and least active MDisk groups in a storage subsystem. To monitor the workload activity of MDisk groups, you can run customized reports such as the MDisk Groups - Workload Activity Report. This formula is used for MDisk groups in the following storage systems:

- SAN Volume Controller
- Storwize® V3500
- Storwize V3700
- Storwize V7000
- Storwize V7000 Unified

The formula that is used by IBM Spectrum Control to estimate the read I/O capability rate is:

$$[\text{Estimated Read I/O Workload} * \text{Type of Disk I/O Capability} * \text{Number of Disks}] / [\text{Estimated Read I/O Workload} * (1 - \text{Cache Hit Ratio}) + \text{Estimated Write I/O Workload} * \text{RAID Type Weighted I/O Ratio}]$$

There are two types of value that are used by the formula:

- Fixed values provided by IBM Spectrum Control
- Variable values that you must specify for each MDisk group

The following fixed values are used by the formula:

Estimated Read I/O Workload

The estimated read I/O workload of an MDisk group. The estimate is set to a fixed value of 0.7 (70%).

Estimated Write I/O Workload

The estimated write I/O workload of an MDisk group. The estimate is set to a fixed value of 0.3 (30%).

You specify the following variable values for each MDisk group:

Type of Disk I/O Capability

When you specify the type of disk that the back-end pool uses:

- You can accept the default I/O capability per second of the disk
- You can enter a new value for the I/O capability per second of the disk

For example, the default value for the I/O capability rate per second of fiber types of disk with 15,000 revolutions per minute (rpm) is 150 I/O operations per second. To change the I/O capability for a type of disk, you issue the **setbackenddisktype** command and you enter a new value for the -iops parameter.

Number of Disks

You specify the number of physical disks that the back-end pool uses.

Cache Hit Ratio

When you specify the type of storage system that the back-end pool uses:

- You can accept the default cache hit ratio for the storage system
- You can enter a new value for the cache hit ratio of the storage system

For example, the default cache hit ratio of the XIV® is 0.7 (70%). To change the cache hit ratio for a type of storage system, you issue the **setbackendtype** command and you enter a new value for the -cachehit parameter.

RAID Type Weighted I/O Ratio

When you specify the type of RAID that the back-end pool uses:

- You can accept the default value for the weighted I/O ratio of the type of RAID
- You can enter a new value for the weighted I/O ratio of the type of RAID

For example the weighted I/O ratio for RAID5 types of RAID is 4:1. To change the weighted I/O ratio for a type of RAID, you issue the **setbackendraidtype** command and you enter a new value for the -weightedio parameter.

Tip: In the IBM Spectrum Control GUI, you can set the values for the back-end properties of an MDisk group. You can also use the IBM Spectrum Control CLI called `tpctool` to set the values for back-end MDisk groups. See the **setarray** command. To change the variable values that are used by the Read I/O capability formula, you must use the IBM Spectrum Control CLI.

Example

The following values have been set for a back-end pool. The type of storage system is the XIV. The type of disk is F15 (Fiber - 15 000 rpm), and the type of RAID is X.

Name of variable	Value
Cache Hit Ratio	0.5 (50%)
Type of Disk I/O Capability	150
RAID Type Weighted I/O Ratio	4
Number of Disks	10

Based on the values that are provided, the read I/O capability per second is:

$[0.7 * 150 * 10] / [0.7 * (1 - 0.5) + 0.3 * 4]$

The read I/O capability per second rate is 677. (The result of the calculation is rounded to the nearest whole number.)

Related reference

- [setarray](#)
- [setbackendtype](#)
- [setbackenddisktype](#)
- [setbackendraidtype](#)

Changing threshold values in storage tier reports

Use the **setdscfg** command to change the threshold values in workload activity reports for managed disk groups and virtual disks.

Before you begin

To change threshold values, you must know the value of the property key that is associated with the name of the threshold in the report.

Procedure

- To start the IBM Spectrum® Control CLI, issue the **tpctool** command.
- Issue the **setdscfg** command, and then set the following values by using the **-property** parameter:
 - Specify the name of the threshold with the *property_key* variable.
 - Specify the new threshold value with the *value* variable.

The following sample code sets the read I/O rate for MDisk groups to 80%:

```
tpctool setdscfg -url localhost:9550 -user ***** -pwd ***** -property  
ManagedDiskGroupAvailIORate 80
```

- Press Enter.

Results

The value of the threshold is changed.

Related reference

- [Threshold values in storage tier reports](#)

Threshold values in storage tier reports

Threshold values are used to assess the workload activity of managed disk (MDisk) groups and virtual disks (VDisks) in storage tier reports. To change a threshold value, you must know the value of the property key that corresponds with the threshold value that is used in the report.

The following table contains the name of the threshold value and the corresponding value of the property key for the following reports:

- The MDisk Groups - Workload Activity Report
- The MDisk Group Details Report

Threshold name in report	Property key value
Actual Size Threshold	<i>ManagedDiskGroupAvailSize</i>
Read I/O Rate Threshold	<i>ManagedDiskGroupAvailIORate</i>
Max Read Response Time Threshold	<i>ManagedDiskGroupBERResponseExceeds</i>
Max Write Response Time Threshold	<i>ManagedDiskGroupBEWResponseExceeds</i>

The following table contains the name of the threshold value and the corresponding value of the property key for the following reports:

- MDisk Groups - VDisk Workload Activity Report
- MDisk Groups - VDisk Workload Activity Details Report
- VDisk Details Report
- Storage Resource Group - VDisk Workload Activity Report

Threshold name in report	Property key value
VDisk Max Read Response Time Threshold	VDiskFERResponseExceeds
VDisk Max Write Response Time Threshold	VDiskFEWResponseExceeds
Max Read Cache Hits Time Threshold	VDiskFERHitExceeds
Max Read I/O Rate Threshold	VDiskFEIORExceeds
Max Write I/O Rate Threshold	VDiskFEIOWExceeds

Related tasks

- [Changing threshold values in storage tier reports](#)

Custom reports about performance

You can use the optional Cognos® Analytics reporting tool to create custom performance reports. Performance reports can contain detailed information about the performance of monitored resources, and some of the properties of those resources.

The resources for which you can create performance reports are storage systems and their components, storage virtualizers and their components, and switches and switch ports. You can add statistics for volumes, back-end arrays, caches, ports, port errors, and other statistics to performance reports. For example, you can add key statistics to a report to create and generate a performance chart for a specific volume of a storage system.

Performance data is collected at intervals. An interval represents the number of minutes over which samples of performance data are averaged. Data that is collected at certain intervals is automatically consolidated, or rolled up, to higher intervals. For example, data collected at 1-minute intervals is consolidated into 5-minute data. When you create a performance report in the Cognos Analytics reporting tool, the smallest interval for the sample data is 5 minutes. The data collected at 1-minute intervals is not available in the Cognos Analytics reporting tool to avoid performance problems.

- [Creating custom performance reports](#)
You can create reports that show performance metrics in the Cognos Analytics reporting tool. Performance reports can contain detailed information about the performance of monitored resources and some of the properties of those resources. You access custom performance reports from IBM Spectrum Control, and you create the reports in the Cognos Analytics reporting tool.
- [Data and properties in performance reports](#)
You can include capacity data, properties, and other information about storage systems, storage virtualizers, and switches in performance reports.

Creating custom performance reports

You can create reports that show performance metrics in the Cognos® Analytics reporting tool. Performance reports can contain detailed information about the performance of monitored resources and some of the properties of those resources. You access custom performance reports from IBM Spectrum® Control, and you create the reports in the Cognos Analytics reporting tool.

Before you begin

When you create a performance report, include properties and metrics for only one type of device in the report.

The list that you use to create performance reports does not include information about the relationships between storage resources. If you include different types of resources in the report, the report might return incorrect results. For example, if you include storage virtualizers and storage systems in a report, the report might return incorrect results.


Tip: Before you run a report, you can preview the report with no data.

About this task

You must specify the time interval for which you want to view performance information. You can specify the time interval in the following ways:

- Select a time unit, for example, day, day of week, or week of month.
- Select performance metrics from sample data, hourly data, or daily data in the resource for which you want performance information. You can also add the time and the interval length to reports. The time is the time at which the performance data was collected. The interval length is the length of the data interval in seconds for sample data, or in minutes for hourly data or daily data.

Procedure

- Go to the URL for your Cognos Analytics server. The format of the URL is similar to this URL: `http://myhostname:9300/bi`
- Depending on the type of report that you want to see, select one of the following items:
 - To create a basic report with limited formatting about the performance of resources, complete the following steps:
 - Click New  in the Welcome portal.
 - Click Other, then click Query Studio.
 - Click IBM Spectrum Control Packages, then click the Performance package.
 - To create a report about the performance of resources and apply advanced formatting features to the report, complete the following steps:
 - In the Welcome portal, click Team Content.

- Click IBM Spectrum Control Packages.
 - Right-click Performance, then click Create report.
 - Click a template for the report.
3. Expand IBM Spectrum Control in the list of data sources.
 4. Expand Performance.
 5. Optional: Select a time unit from the Time item.
 6. Explore the folders in the list to locate the resources, properties, and metrics that you want to include in the report.
 7. Depending on your browser, add properties and metrics to the report:
 - In Internet Explorer, drag the items that you want to see information about to the work area.
 - In Firefox, select the items that you want to see information about, and then click Insert.
 8. Use filters with your performance reports to ensure that your reports do not include too much data. To use filters with your performance reports, do the following steps:
 - a. Add a resource or resource type to the report.
 - b. Select the column heading for the resource or resource type in the work area, and then filter the name of the resource or resource type if necessary.
 - c. Locate the Sample Data, Hourly Data, or Daily Data folder for the resource or resource type in the list. Add a time stamp from the list to the work area. For example, you can add a day time stamp to the report.
 - d. Select the column heading for the time stamp in the work area, and then filter the time.

If you do not enter a time stamp, the report is run on all the records for this resource or resource type in IBM Spectrum Control.
 9. Click Save.

Related information

-  [Run a report](#)

Data and properties in performance reports

You can include capacity data, properties, and other information about storage systems, storage virtualizers, and switches in performance reports.

- [Data for storage systems in performance reports](#)
You can include general information, capacity data, properties, and other information about storage systems in performance reports.
- [Data for storage system volumes in performance reports](#)
You can include general information, capacity data, properties, and other information about storage system volumes in performance reports.
- [Data for storage system ports in performance reports](#)
You can include general information, properties, and other information about storage system ports in performance reports.
- [Data for storage system RAID arrays in performance reports](#)
You can include general information, capacity data, and properties of storage system RAID arrays in performance reports.
- [Data for storage system controllers in performance reports](#)
You can include general information about storage system controllers in performance reports.
- [Data for storage system modules in performance reports](#)
You can include properties of storage system modules, and other information about storage system modules in performance reports.
- [Data for storage system pools in performance reports](#)
You can include general information, capacity data, properties, and other information about storage system pools in performance reports.
- [Data for managed disks on storage systems in performance reports](#)
You can include general information, capacity data, properties, and other information about managed disks on storage systems in performance reports.
- [Data for storage system nodes in performance reports](#)
You can include general information, capacity data, properties, and other information about storage system nodes in performance reports.
- [Data for storage system I/O groups in performance reports](#)
You can include general information, capacity data, properties, and other information about storage system I/O groups in performance reports.
- [Data for local disks on storage systems in performance reports](#)
You can include general information, capacity data, properties, and other information about local disks that are on storage systems in performance reports.
- [Data for storage host connections in performance reports](#)
You can include general information, capacity data, properties, and other information about storage host connections in performance reports.
- [Data for switches in performance reports](#)
You can include properties and other information about switches in performance reports.
- [Data for switch ports in performance reports](#)
You can include properties and other information about switch ports in performance reports.

Data for storage systems in performance reports

You can include general information, capacity data, properties, and other information about storage systems in performance reports.

Information about storage systems

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Capacity and usage data

You can create performance reports that include the following information:

Storage System Pool Capacity (GiB)

The amount of storage space in pools that are on the storage system. For an XIV® or IBM Spectrum Accelerate, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage System Used Pool Space (GiB)

The amount of space that is in use in all pools on a storage system.

Storage System Available Pool Space (GiB)

The amount of unused space in pools that are on the storage system.

Storage System Volume Capacity (GiB)

The amount of space on all volumes on the storage system.

Storage System Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage System Assigned Volume Space (GiB)

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage System Unallocated Disk Space (GiB)

The amount of disk space that can be added to a pool.

Storage System Number of Disks

The number of physical disks on the storage system. For a resource that is running IBM Spectrum Virtualize and is configured as a back-end device, the value is the number of managed disks on the resource.

Storage System Number of Volumes

The number of volumes on a resource.

Storage System Cache (GiB)

The size of the cache on the storage system. This value is not shown for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage System Real Available Pool Space (GiB)

The amount of unused space in pools that have an associated soft size on an XIV or IBM Spectrum Accelerate. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Real Configured Pool Space (GiB)

The amount of storage space that is in pools that have an associated soft size on an XIV or IBM Spectrum Accelerate. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Volume Capacity for z/OS® (GiB)

The amount of space on all volumes on the storage system that the z/OS operating system can use.

Storage System Volume Capacity Assigned to MDisk (GiB)

The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

Storage System Disk Capacity (GiB)

The amount of storage space on physical disks in a storage system, excluding spare disks. This value is not shown for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage System Total Disk Capacity (GiB)

The amount of space on physical disks on a storage system, including spare disks.

Storage System Physical Allocation Percentage

The percentage of physical space in storage system pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{allocated pool space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a 200 GiB storage pool. Therefore, the space that is reserved for volumes is 50 GiB.

Storage System Virtual Allocation Percentage

The percentage of physical space in storage system pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage System Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

Storage System Shortfall Percentage

The percentage of the remaining unallocated volume space in storage system pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{volume space} - \text{used volume space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage System MDisk Available Space (GiB)

The amount of storage space that is available on a managed disk. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage System MDisk Capacity (GiB)

The amount of storage space on the managed disk.

Component properties

You can create performance reports that include the following information:

Storage System Access Information

The URL from which you can access the storage system or storage virtualizer.

Storage System Alias

An alternative name for a storage system. The user defines this name.

Storage System Code Level

The Shared Ethernet Adapter level of a DS8000® storage system. For other storage systems, this value is the firmware version.

Storage System Time Zone

The time zone in which a resource is located.

Storage System Custom Tag 1, 2, and 3

User-defined text that is associated with a storage system. You can add or edit the custom tags for a storage system in the Properties pane of the storage system.

Storage System Location

The physical location of a storage system. The location is defined when a storage system is added to IBM Spectrum Control. You can add or edit the location of the storage system in the Properties pane of the storage system.

Storage System User Provided Name

The name that was specified for the storage system in the storage environment.

Storage System Is Compression Active

Shows whether the compression feature is enabled on the storage system. If this value is **Yes**, the compression feature is enabled.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage System Status

The condition of the resource, for example normal, warning, or error.

Storage System Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Storage System Last Data Collection

The date and time when storage statistics were last collected from the resource.

Storage System Last Data Collection Status

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

Vendor, model, and device information

You can create performance reports that include the following information:

Storage System Type

The type of storage system. For example, the storage system can be an IBM System Storage DS8800 system, an IBM System Storage DS8700 system, an IBM System Storage XIV system, or another type of storage system.

Storage System Vendor

The vendor who supplied the resource.

Storage System Model

The model name or model number of the storage system.

Storage System Machine Type

The machine type of a storage system. For example, the storage system can be DS8800, DS8700, XIV - 2812, or another type of storage system.

Storage System IP Address

The IP address of the resource.

Storage System Serial Number

The serial number of the resource.

- [Performance metrics for SAN Volume Controller and Storwize systems](#)
A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports, port errors, volumes, caches, back-end arrays, and other data for these systems in performance reports.
- [Performance metrics for DS8000 storage systems](#)
You can include performance metrics for ports, port errors, volumes, caches, back-end arrays, and other data for DS8000 storage systems in performance reports.
- [Performance metrics for XIV systems and IBM Spectrum Accelerate](#)
You can include performance data for volumes, caches, and other data for XIV systems and IBM Spectrum Accelerate in performance reports.
- [Performance metrics for block server systems](#)
You can include performance metrics for volumes, and cache data for block server storage systems in performance reports.

Performance metrics for SAN Volume Controller and Storwize systems

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports, port errors, volumes, caches, back-end arrays, and other data for these systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)
The average number of I/O operations per second for operations in which data is sent from a port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive I/O Rate (ops/s)
The average number of I/O operations per second for operations in which the port receives data. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port I/O Rate (ops/s)
The average number of send operations and receive operations per second.

Port to Host Send I/O Rate (ops/s)
The average number of IOs per second that are sent by the storage system to the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Port to Host Receive I/O Rate (ops/s)
The average number of IOs per second that are received by the storage system from the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Total Port to Host I/O Rate (ops/s)
The average number of IOs per second that are transmitted between the storage system and the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Port to Disk Send I/O Rate (ops/s)
The average number of IOs per second that are sent from the storage system to the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.

Port to Disk Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)
The average number of IOs per second that are transmitted between the storage system and the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.

Port to Local Node Send I/O Rate (ops/s)
The average number of IOs per second that are sent to other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Port to Local Node Receive I/O Rate (ops/s)
The average number of IOs per second that are received from other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Total Port to Local Node I/O Rate (ops/s)
The average number of IOs per second that are transmitted between the resource and other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Port to Remote Node Send I/O Rate (ops/s)
The average number of IOs per second that are sent to nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Port to Remote Node Receive I/O Rate (ops/s)
The average number of IOs per second that are received from nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Total Port to Remote Node I/O Rate (ops/s)
The average number of IOs per second that are transmitted between the resource and nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Port Send Data Rate (MiB/s)
The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send Data Rate (MiB/s)
The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)
The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)
The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Sph_port_send_bandwidth_percentageend Data Rate (MiB/s)
The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)
The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)
The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Send Data Rate (MiB/s)
The average rate at which data is sent to other nodes that are in the local cluster. The rate is measured in MiB per second.

Port to Local Node Receive Data Rate (MiB/s)
The average rate at which data is received from other nodes that are in the local cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and other nodes that are in the local cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)
The average rate at which data is sent to nodes that are in the remote cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)
The average rate at which data is received from nodes that are in the remote cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and nodes that are in the remote cluster. The rate is measured in MiB per second.

Port to Local Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to another node that is in the local cluster. This value represents the external response time of the transfers.

Port to Local Node Receive Response Time (ms/op)

The average number of milliseconds to complete a receive operation from another node that is in the local cluster. This value represents the external response time of the transfers.

Overall Port to Local Node Response Time (ms/op)

The average number of milliseconds to complete a send or receive operation with another node that is in the local cluster. This value represents the external response time of the transfers.

Port to Local Node Send Queue Time (ms/op)

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local cluster.

Port to Local Node Receive Queue Time (ms/op)

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local cluster.

Overall Port to Local Node Queue Time (ms/op)

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local cluster.

Port to Remote Node Send Response Time (ms/op)

The average number of milliseconds to complete a send operation to a node that is in the remote cluster. This value represents the external response time of the transfers.

Port to Remote Node Receive Response Time (ms/op)

The average number of milliseconds to complete a receive operation from a node that is in the remote cluster. This value represents the external response time of the transfers.

Overall Port to Remote Node Response Time (ms/op)

The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote cluster. This value represents the external response time of the transfers.

Port to Remote Node Send Queue Time (ms/op)

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote cluster.

Port to Remote Node Receive Queue Time (ms/op)

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote cluster.

Overall Port to Remote Node Queue Time (ms/op)

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote cluster.

Port error data

You can create performance reports that include the following information:

Link Failure Rate (count/s)

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

Loss of Sync Rate (count/s)

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled. Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

CRC Error Rate (count/s)

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.

Primitive Sequence Protocol Error Rate (count/s)

The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)

The average number of bit errors per second that are detected.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)

The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Overflow I/O Rate (ops/s)

The average number of tracks per second that are written but are delayed because there is not enough space in the write cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Flush Through I/O Rate (ops/s)

The average number of tracks per second that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Write Through I/O Rate (ops/s)

The average number of tracks per second that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred per second from the disks to the cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer.

Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Cache to Disk Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred from the cache to the disks. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Delay Percentage

The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Read Ahead Percentage of Cache Hits

The percentage of all read cache hits that occur on pre-staged data. This value applies only to the volume copy cache if the resource is running IBM Spectrum Virtualize 7.3 or later.

Dirty Write Percentage of Cache Hits

The percentage of all cache write hits that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Overflow Percentage

The percentage of write operations that are delayed because there is not enough space in the write cache.

Write Cache Flush Through Percentage

The percentage of tracks that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Write Through Percentage

The percentage of tracks that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation by a node.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)
The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write I/O Rate (op/s)
The average number of overlapping write operations per second that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.
This value applies to resources that are running IBM Spectrum Virtualize.

Peak Read Response Time (ms)
The worst response time measured for a read operation in the sample interval.

Peak Write Response Time (ms)
The worst response time measured for a write operation in the sample interval.

Global Mirror Secondary Write Lag (ms/op)
The average number of additional milliseconds that it takes to service each secondary write operation for Global Mirror. This value does not include the time to service the primary write operations.
You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.
Average number of additional milliseconds it took to service each secondary write operation for Global Mirror, beyond the time needed to service the primary writes

Processor Utilization Percentage
The average percentage of time that the processors on nodes are busy doing system I/O tasks. This value applies only to resources that are running IBM Spectrum Virtualize.

Global Mirror Overlapping Write Percentage
The percentage of overlapping write operations that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.
Applies to resources that are running IBM Spectrum Virtualize.

Overall Host Attributed Response Time Percentage
The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.
For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Zero Buffer-to-Buffer Credit Timer (microseconds)
The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.
In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that metadata was collected.
If this metric is not available, use the Port Send Delay Time metric instead.

Zero Buffer-to-Buffer Credit Percentage
The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that metadata was collected. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

Performance metrics for DS8000 storage systems

You can include performance metrics for ports, port errors, volumes, caches, back-end arrays, and other data for DS8000* storage systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)

The average number of I/O operations per second for operations in which data is sent from a port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive I/O Rate (ops/s)

The average number of I/O operations per second for operations in which the port receives data. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

Port PPRC Send I/O Rate (ops/s)

The average number of operations per second that are sent by using the PPRC protocol.

Port PPRC Receive I/O Rate (ops/s)

The average number of operations per second that are received by using the PPRC protocol.

Total Port PPRC I/O Rate (ops/s)

The average number of send operations and receive operations per second using the PPRC protocol.

Port Send Data Rate (MiB/s)

The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)

The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port PPRC Send Data Rate (MiB/s)

The average number of MiB per second that are sent by using the PPRC protocol.

Port PPRC Receive Data Rate (MiB/s)

The average number of MiB per second that are received by using the Peer-to-Peer Remote Copy (PPRC) protocol.

Total Port PPRC Data Rate (MiB/s)

The average number of MiB per second that are transferred by using the PPRC protocol. This value includes both send and receive PPRC operations.

Port Send Response Time (ms/op)

The average number of milliseconds to complete a send operation.

Port Receive Response Time (ms/op)

The average number of milliseconds to complete a receive operation.

Overall Port Response Time (ms/op)

The average number of milliseconds to complete a send or receive operation.

Port PPRC Send Response Time (ms/op)

The average number of milliseconds to complete a send operation by using the PPRC protocol.

Port PPRC Receive Response Time (ms/op)

The average number of milliseconds to complete a receive operation by using the PPRC protocol.

Overall Port PPRC Response Time (ms/op)

The average number of milliseconds to complete a send or receive operation by using the PPRC protocol.

Port Send Transfer Size (KiB/op)

The average number of KiB that are transferred per send operation. A send operation is a read operation that is processed or a write operation that is initiated by the port.

Port Receive Transfer Size (KiB/op)

The average number of KiB that are transferred per receive operation. A receive operation is a write operation that is processed or a read operation that is initiated by the port.

Overall Port Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Port error data

You can create performance reports that include the following information:

Link Failure Rate (count/s)

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

Loss of Sync Rate (count/s)

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled. Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

CRC Error Rate (count/s)

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.

Primitive Sequence Protocol Error Rate (count/s)

The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)

The average number of bit errors per second that are detected.

Error Frame Rate (count/s)
The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

Link Reset Transmitted Rate (count/s)
The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

Link Reset Received Rate (count/s)
The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

Out of Order Data Rate (count/s)
The average number of times per second that an out-of-order frame is detected.

Out of Order ACK Rate (count/s)
The average number of times per second that an out-of-order acknowledge (ACK) frame is detected.
An ACK frame is used for end-to-end flow control and is sent to verify receipt of a frame.

Duplicate Frame Rate (count/s)
The average number of duplicate frames per second that are received. A duplicate frame is a frame that the system previously processed for the port.

Invalid Relative Offset Rate (count/s)
The average number of times per second that frames are received with an invalid relative offset in the frame header.

Sequence Timeout Rate (count/s)
The average number of times per second that the port detects a timeout condition after the port receives a sequence initiative for a Fibre Channel exchange.

Credit Recovery Link Reset Rate
The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

Volume data

You can create performance reports that include the following information:

Normal Read I/O Rate (ops/s)
The average number of nonsequential read operations per second.

Sequential Read I/O Rate (ops/s)
The average number of sequential read operations per second.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Normal Write I/O Rate (ops/s)
The average number of nonsequential write operations per second.

Sequential Write I/O Rate (ops/s)
The average number of sequential write operations per second.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Normal I/O Rate (ops/s)
The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Sequential I/O Rate (ops/s)
The average number of sequential I/O operations per second. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Record Mode Read I/O Rate (ops/s)
The average number of I/O operations per second for record-mode read operations.
For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

Read Data Rate (MiB/s)
The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Normal Read Cache Hit Percentage
The percentage of nonsequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Sequential Read Cache Hit Percentage
The percentage of sequential read operations that find data in the cache. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Normal Write Cache Hit Percentage
The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage
The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Normal Cache Hit Percentage
The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage
The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)
The average number of sectors or tracks per second that are transferred per second from the disks to the cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Cache to Disk Transfer Rate (ops/s)
The average number of sectors or tracks per second that are transferred from the cache to the disks. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Record Mode Read Cache Hit Percentage
The percentage of cache hits for record-mode read operations. For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)

The average number of read operations per second that are issued by the High Performance FICON® feature of the storage system.

HPF Write I/O Rate (ops/s)

The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

Cache Holding Time (s)

The average number of seconds that I/O data for a storage system node is held in the cache. A short cache-holding time indicates adverse performance.

PPRC Transfer Rate (ops/s)

The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Performance metrics for XIV systems and IBM Spectrum Accelerate

You can include performance data for volumes, caches, and other data for XIV® systems and IBM Spectrum Accelerate in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Read Data Cache Hit Percentage

The percentage of all read data that is read from the cache.

Write Data Cache Hit Percentage

The percentage of all write data that is written to cache slots that are marked as modified.

Overall Data Cache Hit Percentage

The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

SSD Read Cache Hit Percentage

The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

SSD Read Data Cache Hit Percentage

The percentage of all read data that was read from cache memory on a solid-state drive.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

Read Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a read-cache hit operation.

Write Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a write-cache hit operation.

Overall Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

Read Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a read-cache miss operation.

Write Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a write-cache miss operation.

Overall Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

SSD Read Cache Hit Response Time (ms/op)

The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

Other data

You can create performance reports that include the following information:

Small Transfers I/O Percentage

The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers I/O Percentage

The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers I/O Percentage

The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers I/O Percentage

The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

Performance metrics for block server systems

You can include performance metrics for volumes, and cache data for block server storage systems in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Data for storage system volumes in performance reports

You can include general information, capacity data, properties, and other information about storage system volumes in performance reports.

Information about storage system volumes

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Controller Name

The name that was assigned to the controller when it was added to the system.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage RAID Array Name

The name that was assigned to the RAID array when it was added to the system.

Storage I/O Group Name

The name that was assigned to the I/O group when it was added to the system.

Storage Preferred Node Name

The name of the storage node that is used for I/O operations for a volume.

Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

Storage Volume Assigned Host Connection

The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

Storage Volume Host Connection OS Type

The operating system type of the server or hypervisor that the volume is assigned to.

Storage System Volume Group ID

The unique identifier for a collection of volumes on DS8000® storage systems. The storage system generates this identifier.

Storage Volume Group Name

The user-defined name for a collection of volumes on DS8000 storage systems.

Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Hypervisor Name

The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be `hypervisor.example.com`.

Hypervisor Cluster Name

The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

Capacity and usage data

You can create performance reports that include the following information:

Storage Volume Capacity (GiB)

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV® systems and IBM Spectrum Accelerate, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Volume Used Space (GiB)

The amount of allocated space that is used by a volume.

For resources that are running IBM Spectrum Virtualize, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that IBM Spectrum Control collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Volume Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Volume Allocated Space (GiB)

The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following resources:

- Resources other than those that are running IBM Spectrum Virtualize and are configured as back-end storage
- Resources that are running IBM Spectrum Virtualize, are configured as back-end storage, and are not thin-provisioned

Storage Volume Unallocated Space (GiB)

The amount of space in a pool that is not reserved for a volume.

IBM Spectrum Control uses the following formula to determine this value:

capacity - allocated space

Available only for thin provisioned volumes.

Storage Volume Physical Allocation Percentage

The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

IBM Spectrum Control uses the following formula to determine this value:

(allocated space ÷ capacity) × 100

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Volume Shortfall Percentage

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume.

The higher the percentage, the more critical the shortfall of space.

IBM Spectrum Control uses the following formula to determine this value:

(unallocatable volume space ÷ volume unallocated space) × 100

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This Storage Volume Shortfall Percentage represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

**100 GiB unallocated volume space - 70 GiB remaining
pool space ÷ 100 GiB unallocated volume space × 100**

Storage Volume Unallocatable Space (GiB)

The unallocatable storage space in GiB of a thin-provisioned volume. Unallocatable space cannot be supplied by the pool to which the volume belongs.

Storage Volume Used Allocated Space Percentage

The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

IBM Spectrum Control uses the following formula to determine this value:

(volume used space ÷ volume allocated space) × 100

This property is available only for volumes on resources that are running IBM Spectrum Virtualize.

Storage Volume Unused Space (GiB)

The amount of space that is allocated to a volume and is not yet used.

IBM Spectrum Control uses the following formula to determine this value:

allocated space - used space

This property is available only for pools on resources that are running IBM Spectrum Virtualize.

Storage Volume Uncompressed Capacity (GiB)

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Component properties

You can create performance reports that include the following information:

Storage Volume RAID Level

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Volume Is Thin Provisioned

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is **Yes**, the resource is thin-provisioned.

Storage Volume Is Assigned
Shows whether the volume is assigned to a server. If this value is **Yes**, the volume is assigned to a server.

Storage Volume Is Encrypted
Shows whether the resource is encrypted. If this value is **Yes**, the resource is encrypted.

Storage Volume Is Encryptable
Shows whether the resource can be encrypted. If this value is **Yes**, the resource can be encrypted.

Storage Volume Is Volume Control Manager Database
Shows whether a volume is used as the Volume Control Manager Database. If this value is **Yes**, the volume is used as the Volume Control Manager Database.

Storage Volume Format
Shows the format of the volume or pool. The format can be a Count Key Data (CKD) format, or a fixed block format.

Storage Volume Serial Number
The serial number or volume ID of the volume.

Storage Volume Number
The number of the volume that is assigned by the system.

Storage Volume WWN
The worldwide name of the volume.

Storage Volume FlashCopy® Relationship
Shows whether a volume on a storage system is in a FlashCopy relationship. This property can contain the following values:

- Source Relationship**
The volume is the source of the relationship.
- Target Relationship**
The volume is the target of the relationship.
- Both Source and Target Relationship**
The volume is both the source and target of the relationship.
- Not in a FlashCopy Relationship**
The volume is not in a FlashCopy relationship.

Storage Volume Remote Copy Relationship
Shows whether a volume on a storage system is in a remote copy relationship. This property can contain the following values:

- Not in a Copy Relationship**
The volume is not in a remote copy relationship.
- Source Relationship**
The volume is the source of the relationship.
- Target Relationship**
The volume is the target of the relationship.
- Both Source and Target Relationship**
The volume is both the source and target of the relationship.

Storage Volume Logical Subsystem
The logical subsystem (LSS) to which a volume or pool belongs.

Storage Volume Type
The type of storage volume. For example, the storage volume can be striped or sequential. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Fast Write State
Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted.

Storage Volume Grain Size (KiB)
The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.
This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Is Auto Expand Enabled
Shows whether a volume can automatically allocate new extents from a pool. Volumes might allocate new extents to expand the real capacity of the volume. If this value is **Yes**, the volume can automatically allocate new extents from a pool.
This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Metro Mirror Name
The name of the Metro Mirror that keeps the synchronous copy of the storage volume. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Mirror Count
The number of mirrors that keep a synchronous copy of the resource. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Throttle
The maximum number of commands that the volume can queue. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Warning Level
The percentage of volume capacity that is used at which a warning is generated. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Is Compressed
Shows whether the volume is compressed. If this value is **Yes**, the volume is compressed.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage Volume Status

The condition of the resource, for example normal, warning, or error.

Storage Volume Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Storage Volume Native Status

The condition of the volume. The status can be online, offline, degraded, or excluded.

- [Performance metrics for volumes on SAN Volume Controller and Storwize systems](#)

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for volumes, caches, and other data for volumes on these systems in performance reports.

- [Performance metrics for volumes on DS8000 storage systems](#)

You can include performance data for volumes, caches, and other data for storage system volumes on DS8000 storage systems in performance reports.

- [Performance metrics for storage system volumes on XIV systems and IBM Spectrum Accelerate](#)

You can include performance data for volumes, caches, and other data for storage system volumes on XIV systems and IBM Spectrum Accelerate in performance reports.

- [Performance metrics for storage system volumes on block servers](#)

You can include performance data for volumes and caches for storage system volumes on block servers in performance reports.

Performance metrics for volumes on SAN Volume Controller and Storwize systems

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for volumes, caches, and other data for volumes on these systems in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Maximum Read I/O Rate (ops/s)

The maximum read I/O rate of the volume in operations per second. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Maximum Write I/O Rate (ops/s)

The maximum write I/O rate of the volume in operations per second. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Maximum Write I/O Size (KiB)

The maximum size of write I/O operations in kibibytes. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Maximum Read Response Time (ms/op)

The maximum read response time of the volume in milliseconds per operation. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Maximum Write Response Time (ms/op)

The maximum write response time of the volume in milliseconds per operation. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Volume Utilization

The average percentage of time that the volume is busy.

Maximum Volume Utilization Percentage

The maximum period in percentage that the volume is busy. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Cache data

You can create performance reports that include the following information:

Write Cache Overflow I/O Rate (ops/s)

The average number of tracks per second that are written but are delayed because there is not enough space in the write cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Delay I/O Rate (ops/s)

The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Flush Through I/O Rate (ops/s)

The average number of tracks per second that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Write Through I/O Rate (ops/s)

The average number of tracks per second that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Maximum Write Cache Delay I/O Rate (ops/s)

The maximum number of delayed I/O operations per second for the volume, where the delay occurred because of space constraints or other conditions.

This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred per second from the disks to the cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Cache to Disk Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred from the cache to the disks. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Delay Percentage

The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Read Ahead Percentage of Cache Hits

The percentage of all read cache hits that occur on pre-staged data. This value applies only to the volume copy cache if the resource is running IBM Spectrum Virtualize 7.3 or later.

Dirty Write Percentage of Cache Hits

The percentage of all cache write hits that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Overflow Percentage

The percentage of write operations that are delayed because there is not enough space in the write cache.

Write Cache Flush Through Percentage

The percentage of tracks that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Write Through Percentage

The percentage of tracks that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Maximum Read Cache Hit Percentage

The maximum number of cache hits in percentage for sequential and nonsequential read operations for the volume over a specified interval. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)

The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write Percentage

The percentage of overlapping write operations that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.

Applies to resources that are running IBM Spectrum Virtualize.

Global Mirror Overlapping Write I/O Rate (op/s)

The average number of overlapping write operations per second that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.

This value applies to resources that are running IBM Spectrum Virtualize.

Peak Read Response Time (ms)

The worst response time measured for a read operation in the sample interval.

Peak Write Response Time (ms)

The worst response time measured for a write operation in the sample interval.

Global Mirror Secondary Write Lag (ms/op)

The average number of additional milliseconds that it takes to service each secondary write operation for Global Mirror. This value does not include the time to service the primary write operations.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Average number of additional milliseconds it took to service each secondary write operation for Global Mirror, beyond the time needed to service the primary writes

Overall Host Attributed Response Time Percentage

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Nonpreferred Node Usage Percentage

The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.

Performance metrics for volumes on DS8000 storage systems

You can include performance data for volumes, caches, and other data for storage system volumes on DS8000® storage systems in performance reports.

Volume data

You can create performance reports that include the following information:

Normal Read I/O Rate (ops/s)

The average number of nonsequential read operations per second.

Sequential Read I/O Rate (ops/s)

The average number of sequential read operations per second.

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Normal Write I/O Rate (ops/s)

The average number of nonsequential write operations per second.

Sequential Write I/O Rate (ops/s)

The average number of sequential write operations per second.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Normal I/O Rate (ops/s)

The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Sequential I/O Rate (ops/s)

The average number of sequential I/O operations per second. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Record Mode Read I/O Rate (ops/s)

The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Volume Utilization

The average percentage of time that the volume is busy.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)

The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Normal Read Cache Hit Percentage

The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Sequential Read Cache Hit Percentage

The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Normal Write Cache Hit Percentage

The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage

The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Normal Cache Hit Percentage

The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage

The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred per second from the disks to the cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Cache to Disk Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred from the cache to the disks. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Delay Percentage

The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Record Mode Read Cache Hit Percentage

The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)

The average number of read operations per second that are issued by the High Performance FICON® feature of the storage system.

HPF Write I/O Rate (ops/s)

The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

PPRC Transfer Rate (ops/s)

The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Performance metrics for storage system volumes on XIV systems and IBM Spectrum Accelerate

You can include performance data for volumes, caches, and other data for storage system volumes on XIV® systems and IBM Spectrum Accelerate in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)
The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Volume Utilization
The average percentage of time that the volume is busy.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.
You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.
You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Read Data Cache Hit Percentage
The percentage of all read data that is read from the cache.

Write Data Cache Hit Percentage
The percentage of all write data that is written to cache slots that are marked as modified.

Overall Data Cache Hit Percentage
The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

Read Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a read-cache miss operation.

Write Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a write-cache miss operation.

Overall Cache Miss Response Time (ms/op)
The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

Read Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a read-cache hit operation.

Write Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a write-cache hit operation.

Overall Cache Hit Response Time (ms/op)
The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

SSD Read Cache Hit Percentage
The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric.
This metric is only available for XIV systems that use solid-state drives as drives for caching.

SSD Read Data Cache Hit Percentage
The percentage of all read data that was read from cache memory on a solid-state drive.
This metric is only available for XIV systems that use solid-state drives as drives for caching.

SSD Read Cache Hit Response Time (ms/op)
The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric.
This metric is only available for XIV systems that use solid-state drives as drives for caching.

Other data

You can create performance reports that include the following information:

Small Transfers I/O Percentage
The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers I/O Percentage
The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers I/O Percentage

The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers I/O Percentage

The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

Performance metrics for storage system volumes on block servers

You can include performance data for volumes and caches for storage system volumes on block servers in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Volume Utilization

The average percentage of time that the volume is busy.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Data for storage system ports in performance reports

You can include general information, properties, and other information about storage system ports in performance reports.

Information about storage system ports

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Module Name

The name that was assigned to the module when it was added to the system.

Storage Port Name

The name that was assigned to the storage port when the storage system was added to the system.

Component properties

You can create performance reports that include the following information:

Storage System Port Number

The port number on the resource.

Storage System Port Type

The type of port on the storage system, storage virtualizer, or switch. For example, the port type can be N_Port, F_Port, or another type of port.

Storage System Port Speed (GiB/s)

The speed of a port, which is measured in GiB per second.

Storage System Port Location

The physical location of a port on DS8000® storage systems and XIV® systems. For all other systems, this property shows the WWPN of the port.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage System Port Status

The condition of the resource, for example normal, warning, or error.

Storage System Port Enabled State

Shows whether a port is enabled, disabled, or is enabled but offline.

Storage System Port Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

- [Performance metrics for ports on SAN Volume Controller and Storwize systems](#)

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports and port errors on these systems in performance reports.

- [Performance metrics for ports on DS8000 and DS6000 storage systems](#)

You can include performance data for storage system port and port errors on DS8000 and DS6000™ storage systems in performance reports.

- [Performance metrics for ports on an XIV](#)

You can include performance metrics for storage system ports on an XIV in performance reports.

- [Performance metrics for storage system ports on block servers](#)

You can include performance metrics for storage system ports on block servers in performance reports.

Performance metrics for ports on SAN Volume Controller and Storwize systems

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports and port errors on these systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)

The average number of I/O operations per second for operations in which data is sent from a port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive I/O Rate (ops/s)

The average number of I/O operations per second for operations in which the port receives data. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

Port to Host Send I/O Rate (ops/s)

The average number of IOs per second that are sent by the storage system to the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Port to Host Receive I/O Rate (ops/s)

The average number of IOs per second that are received by the storage system from the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Total Port to Host I/O Rate (ops/s)

The average number of IOs per second that are transmitted between the storage system and the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Port to Disk Send I/O Rate (ops/s)

The average number of IOs per second that are sent from the storage system to the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.

Port to Disk Receive I/O Rate (ops/s)

The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)

The average number of IOs per second that are transmitted between the storage system and the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.

Port to Local Node Send I/O Rate (ops/s)

The average number of IOs per second that are sent to other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Port to Local Node Receive I/O Rate (ops/s)

The average number of IOs per second that are received from other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Total Port to Local Node I/O Rate (ops/s)

The average number of IOs per second that are transmitted between the resource and other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Port to Remote Node Send I/O Rate (ops/s)

The average number of IOs per second that are sent to nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Port to Remote Node Receive I/O Rate (ops/s)

The average number of IOs per second that are received from nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Total Port to Remote Node I/O Rate (ops/s)

The average number of IOs per second that are transmitted between the resource and nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Port Send Data Rate (MiB/s)

The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)

The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send Data Rate (MiB/s)

The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)

The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)

The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Sph_port_send_bandwidth_percentageend Data Rate (MiB/s)

The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)

The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)

The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Send Data Rate (MiB/s)

The average rate at which data is sent to other nodes that are in the local cluster. The rate is measured in MiB per second.

Port to Local Node Receive Data Rate (MiB/s)

The average rate at which data is received from other nodes that are in the local cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)

The average rate at which data is transmitted between the component and other nodes that are in the local cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)

The average rate at which data is sent to nodes that are in the remote cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)

The average rate at which data is received from nodes that are in the remote cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)

The average rate at which data is transmitted between the component and nodes that are in the remote cluster. The rate is measured in MiB per second.

Port Send Bandwidth Percentage

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage

The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port error data

You can create performance reports that include the following information:

Link Failure Rate (count/s)

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

Loss of Sync Rate (count/s)

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled. Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

CRC Error Rate (count/s)

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.

Primitive Sequence Protocol Error Rate (count/s)

The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)

The average number of bit errors per second that are detected.

Zero Buffer-to-Buffer Credit Timer (microseconds)

The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that metadata was collected.

If this metric is not available, use the Port Send Delay Time metric instead.

Zero Buffer-to-Buffer Credit Percentage

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that metadata was collected. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

Port Congestion Index

The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

Performance metrics for ports on DS8000 and DS6000 storage systems

You can include performance data for storage system port and port errors on DS8000® and DS6000™ storage systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)

The average number of I/O operations per second for operations in which data is sent from a port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive I/O Rate (ops/s)

The average number of I/O operations per second for operations in which the port receives data. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

Port FICON Send I/O Rate (ops/s)

The average number of send operations per second for Fibre Channel connection (FICON®) usage.

Port FICON Receive I/O Rate (ops/s)

The average number of receive operations per second for FICON usage.

Total Port FICON I/O Rate (ops/s)

The average number of send and receive operations per second for FICON usage.

Port FCP Send I/O Rate (ops/s)

The average number of send operations per second for Fibre Channel Protocol (FCP) usage.

Port FCP Receive I/O Rate (ops/s)

The average number of receive operations per second for FCP usage.

Total Port FCP I/O Rate (ops/s)

The average number of send operations and receive operations per second for FCP usage.

Port PPRC Send I/O Rate (ops/s)

The average number of operations per second that are sent by using the PPRC protocol.

Port PPRC Receive I/O Rate (ops/s)

The average number of operations per second that are received by using the PPRC protocol.

Total Port PPRC I/O Rate (ops/s)

The average number of send operations and receive operations per second using the PPRC protocol.

Port Send Data Rate (MiB/s)

The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port FICON Send Data Rate (MiB/s)
The average number of MiB per second that is sent for FICON usage.

Port FICON Receive Data Rate (MiB/s)
The average number of MiB per second that is received for FICON usage.

Total Port FICON Data Rate (MiB/s)
The average number of MiB per second that is transferred for FICON usage. This value includes both send and receive FICON operations.

Port FCP Send Data Rate (MiB/s)
The average number of MiB per second that are sent for FCP usage.

Port FCP Receive Data Rate (MiB/s)
The average number of MiB per second that are received for FCP usage.

Total Port FCP Data Rate (MiB/s)
The average number of MiB per second that are transferred for FCP usage. This value includes both send and receive FCP operations.

Port PPRC Send Data Rate (MiB/s)
The average number of MiB per second that are sent by using the PPRC protocol.

Port PPRC Receive Data Rate (MiB/s)
The average number of MiB per second that are received by using the Peer-to-Peer Remote Copy (PPRC) protocol.

Total Port PPRC Data Rate (MiB/s)
The average number of MiB per second that are transferred by using the PPRC protocol. This value includes both send and receive PPRC operations.

Port Send Response Time (ms/op)
The average number of milliseconds to complete a send operation.

Port Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation.

Overall Port Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation.

Port FCP Send Response Time (ms/op)
The average number of milliseconds to complete a send operation for FCP usage.

Port FCP Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation for FCP usage.

Overall Port FCP Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation for FCP usage.

Port FICON Send Response Time (ms/op)
The average number of milliseconds to complete a send operation for FICON usage.

Port FICON Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation for FICON usage.

Overall Port FICON Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation for FICON usage. This value includes both send and receive FICON operations.

Port PPRC Send Response Time (ms/op)
The average number of milliseconds to complete a send operation by using the PPRC protocol.

Port PPRC Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation by using the PPRC protocol.

Overall Port PPRC Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation by using the PPRC protocol.

Port Send Transfer Size (KiB/op)
The average number of KiB that are transferred per send operation. A send operation is a read operation that is processed or a write operation that is initiated by the port.

Port Receive Transfer Size (KiB/op)
The average number of KiB that are transferred per receive operation. A receive operation is a write operation that is processed or a read operation that is initiated by the port.

Overall Port Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Port Send Utilization Percentage
The average percentage of time that the port is busy sending data.

Port Receive Utilization Percentage
The average percentage of time that the port is busy receiving data.

Overall Port Utilization Percentage
The average percentage of time that the port is busy sending or receiving data.

Port Send Bandwidth Percentage
The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage
The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage
The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port error data

You can create performance reports that include the following information:

Error Frame Rate (count/s)

The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

Link Failure Rate (count/s)

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

Loss of Sync Rate (count/s)

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

CRC Error Rate (count/s)

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.

Primitive Sequence Protocol Error Rate (count/s)

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)

The average number of bit errors per second that are detected.

Out of Order Data Rate (count/s)

The average number of times per second that an out-of-order frame is detected.

Out of Order ACK Rate (count/s)

The average number of times per second that an out-of-order acknowledge (ACK) frame is detected.

An ACK frame is used for end-to-end flow control and is sent to verify receipt of a frame.

Duplicate Frame Rate (count/s)

The average number of duplicate frames per second that are received. A duplicate frame is a frame that the system previously processed for the port.

Invalid Relative Offset Rate (count/s)

The average number of times per second that frames are received with an invalid relative offset in the frame header.

Sequence Timeout Rate (count/s)

The average number of times per second that the port detects a timeout condition after the port receives a sequence initiative for a Fibre Channel exchange.

Link Reset Transmitted Rate (count/s)

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

Link Reset Received Rate (count/s)

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

Credit Recovery Link Reset Rate

The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

Performance metrics for ports on an XIV

You can include performance metrics for storage system ports on an XIV® in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)

The average number of I/O operations per second for operations in which data is sent from a port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive I/O Rate (ops/s)

The average number of I/O operations per second for operations in which the port receives data. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

Port Send Data Rate (MiB/s)

The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)

The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Response Time (ms/op)

The average number of milliseconds to complete a send operation.

Port Receive Response Time (ms/op)

The average number of milliseconds to complete a receive operation.

Overall Port Response Time (ms/op)

The average number of milliseconds to complete a send or receive operation.

Port Send Bandwidth Percentage

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage

The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Performance metrics for storage system ports on block servers

You can include performance metrics for storage system ports on block servers in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)

The average number of I/O operations per second for operations in which data is sent from a port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive I/O Rate (ops/s)

The average number of I/O operations per second for operations in which the port receives data. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

Port Send Data Rate (MiB/s)

The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)

The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Response Time (ms/op)

The average number of milliseconds to complete a send operation.

Port Receive Response Time (ms/op)

The average number of milliseconds to complete a receive operation.

Overall Port Response Time (ms/op)

The average number of milliseconds to complete a send or receive operation.

Port Send Transfer Size (KiB/op)

The average number of KiB that are transferred per send operation. A send operation is a read operation that is processed or a write operation that is initiated by the port.

Port Receive Transfer Size (KiB/op)

The average number of KiB that are transferred per receive operation. A receive operation is a write operation that is processed or a read operation that is initiated by the port.

Overall Port Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both send and receive operations.

Data for storage system RAID arrays in performance reports

You can include general information, capacity data, and properties of storage system RAID arrays in performance reports.

Information about storage system RAID arrays

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Controller Name

The name that was assigned to the controller when it was added to the system.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage RAID Array Name

The name that was assigned to the RAID array when it was added to the system.

Capacity and usage data

You can create performance reports that include the following information:

Storage RAID Array DDM Capacity (GiB)

The amount of storage space on the disk drive modules in an array.

Component properties

You can create performance reports that include the following information:

Storage RAID Array RAID Level

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage RAID Array Width

The number of disks that are in the storage array.

- Storage RAID Array DDM Speed (RPM)
The speed of the disk drive module in the array.
- Storage RAID Array Description
Shows the array site that the array is associated with.
- Storage RAID Array Tag
A number that identifies an array.
- Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.
- [Performance metrics for storage system arrays](#)
You can include front end data, cache data, back-end data, and other data for storage system arrays in performance reports.

Performance metrics for storage system arrays

You can include front end data, cache data, back-end data, and other data for storage system arrays in performance reports.

Front end data

You can create performance reports that include the following information:

- Normal Read I/O Rate (ops/s)
The average number of nonsequential read operations per second.
- Sequential Read I/O Rate (ops/s)
The average number of sequential read operations per second.
- Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.
- Normal Write I/O Rate (ops/s)
The average number of nonsequential write operations per second.
- Sequential Write I/O Rate (ops/s)
The average number of sequential write operations per second.
- Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.
- Total Normal I/O Rate (ops/s)
The average number of nonsequential I/O operations per second. This value includes both read and write operations.
- Total Sequential I/O Rate (ops/s)
The average number of sequential I/O operations per second. This value includes both read and write operations.
- Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.
- Record Mode Read I/O Rate (ops/s)
The average number of I/O operations per second for record-mode read operations.
For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.
- Read Data Rate (MiB/s)
The average number of MiBs per second that are transferred for read operations.
- Write Data Rate (MiB/s)
The average number of MiBs per second that are transferred for write operations.
- Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.
- Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.
- Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.
- Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.
- Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.
- Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.
- Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

- Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.
This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.
- Normal Read Cache Hit Percentage
The percentage of nonsequential read operations that find data in the cache.
You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.
- Sequential Read Cache Hit Percentage
The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Normal Write Cache Hit Percentage

The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage

The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Normal Cache Hit Percentage

The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage

The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Record Mode Read Cache Hit Percentage

The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Disk to Cache Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred per second from the disks to the cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer.

Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Cache to Disk Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred from the cache to the disks. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Delay Percentage

The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Back-end data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Sequential Back-End I/O Percentage

The percentage of all I/O operations that are issued to the back-end storage resources. The operations are for the array and are sequential operations.

Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)

The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Back-End Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)

The average number of read operations per second that are issued by the High Performance FICON® feature of the storage system.

HPF Write I/O Rate (ops/s)

The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

PPRC Transfer Rate (ops/s)

The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Disk Utilization Percentage

The average percentage of time that the disks that are associated with an array are busy.

No value is calculated for this property if there are multiple ranks in the extent pool where the thin-provisioned volumes are allocated. In this case, the value **N/A** is displayed. This limitation applies only to DS8000® storage systems.

If there is only a single rank in the extent pool, the value for this property is calculated regardless of the thin-provisioned volumes.

Tip: Some highly sequential workloads such as batch or backup processing might continually exceed the threshold because they drive the arrays to high utilization percentages. For these types of workloads, a high utilization indicates that the work is being performed very efficiently and is not a cause for concern.

For DS8000 storage systems, this property is available only for 8.5.0 and later. For earlier versions, the value **N/A** is shown.

Data for storage system controllers in performance reports

You can include general information about storage system controllers in performance reports.

Information about storage system controllers

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Controller Name

The name that was assigned to the controller when it was added to the system.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

- [Performance metrics for storage system controllers on DS8000 storage systems](#)

You can include front end data, cache data, back-end data, and other data for storage system controllers on DS8000® series systems in performance reports.

- [Performance metrics for storage system controllers on block servers](#)

You can include performance data for volumes and caches for storage system controllers on block servers in performance reports.

Performance metrics for storage system controllers on DS8000 storage systems

You can include front end data, cache data, back-end data, and other data for storage system controllers on DS8000® series systems in performance reports.

Front end data

You can create performance reports that include the following information:

Normal Read I/O Rate (ops/s)

The average number of nonsequential read operations per second.

Sequential Read I/O Rate (ops/s)

The average number of sequential read operations per second.

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Normal Write I/O Rate (ops/s)

The average number of nonsequential write operations per second.

Sequential Write I/O Rate (ops/s)

The average number of sequential write operations per second.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Normal I/O Rate (ops/s)

The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Sequential I/O Rate (ops/s)

The average number of sequential I/O operations per second. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Record Mode Read I/O Rate (ops/s)

The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)

The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Normal Read Cache Hit Percentage

The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Sequential Read Cache Hit Percentage

The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Normal Write Cache Hit Percentage

The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage

The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Normal Cache Hit Percentage

The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage

The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Record Mode Read Cache Hit Percentage

The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Write Cache Delay Percentage

The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Disk to Cache Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred per second from the disks to the cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer.

Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Cache to Disk Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred from the cache to the disks. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Back-end data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)

The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Back-End Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)

The average number of read operations per second that are issued by the High Performance FICON® feature of the storage system.

HPF Write I/O Rate (ops/s)

The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

PPRC Transfer Rate (ops/s)

The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Cache Holding Time (s)

The average number of seconds that I/O data for a storage system node is held in the cache. A short cache-holding time indicates adverse performance.

Performance metrics for storage system controllers on block servers

You can include performance data for volumes and caches for storage system controllers on block servers in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Data for storage system modules in performance reports

You can include properties of storage system modules, and other information about storage system modules in performance reports.

Information about storage system modules

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Module Name

The name that was assigned to the module when it was added to the system.

Component properties

You can create performance reports that include the following information:

Storage Module IP Address

The IP address of the resource.

Storage Module WWN

The worldwide name of the module.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage Module Status

The condition of the resource, for example normal, warning, or error.

Storage Module Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

- [Performance metrics for storage system modules on XIV systems and IBM Spectrum Accelerate](#)

You can include performance metrics for volumes, caches, and other data for storage system modules on XIV® systems and IBM Spectrum Accelerate in performance reports.

Performance metrics for storage system modules on XIV systems and IBM Spectrum Accelerate

You can include performance metrics for volumes, caches, and other data for storage system modules on XIV® systems and IBM Spectrum Accelerate in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Read Data Cache Hit Percentage

The percentage of all read data that is read from the cache.

Write Data Cache Hit Percentage

The percentage of all write data that is written to cache slots that are marked as modified.

Overall Data Cache Hit Percentage

The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

Read Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a read-cache hit operation.

Write Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a write-cache hit operation.

Overall Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

Read Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a read-cache miss operation.

Write Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a write-cache miss operation.

Overall Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

SSD Read Cache Hit Percentage

The percentage of read operations that find data in the cache on a solid-state drive (SSD). The value for this metric is also included in the value for the Overall Read Cache Hit Percentage metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

SSD Read Data Cache Hit Percentage

The percentage of all read data that was read from cache memory on a solid-state drive.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

SSD Read Cache Hit Response Time (ms/op)

The average number of milliseconds that it takes to complete a hit operation on the read cache on a solid-state drive. The value for this metric is also included in the value for the Read Cache Hit Response Time metric.

This metric is only available for XIV systems that use solid-state drives as drives for caching.

Other data

You can create performance reports that include the following information:

Small Transfers I/O Percentage

The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers I/O Percentage

The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers I/O Percentage

The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers I/O Percentage

The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

Data for storage system pools in performance reports

You can include general information, capacity data, properties, and other information about storage system pools in performance reports.

Information about storage system pools

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Capacity and usage data

You can create performance reports that include the following information:

Storage Pool Capacity (GiB)

The total amount of storage space in a pool.

Storage Pool Available Space (GiB)

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

IBM Spectrum Control uses the following formula to determine this value:

pool capacity - used space

For XIV® systems and IBM Spectrum Accelerate, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Storage Pool Total Volume Capacity (GiB)

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Pool Allocated Space (GiB)

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following resources:

- Resources other than those that are running IBM Spectrum Virtualize and are configured as back-end storage
- Resources that are running IBM Spectrum Virtualize, are configured as back-end storage, and are not thin-provisioned

Storage Pool Used Volume Space (GiB)

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For resources that are running IBM Spectrum Virtualize, you can preallocate thin-provisioned volume space when volumes are created. For these resources, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that IBM Spectrum Control collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Pool Extent Size (MiB)

The size of the extent that was specified when a pool was created. Smaller extent sizes limit the maximum size of the volumes that can be created in a pool. Smaller extent sizes minimize the amount of potentially wasted space per volume.

This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Pool Assigned Volume Space (GiB)

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Pool Unassigned Volume Space (GiB)

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Pool Number of Storage Volumes

The number of volumes in the storage pool.

Storage Pool Number of MDisks

The number of managed disks in a storage pool.

Storage Pool Physical Allocation Percentage

The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

IBM Spectrum Control uses the following formula to determine the allocation percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Pool Virtual Allocation Percentage

The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

IBM Spectrum Control uses the following formula to determine the allocation percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes. An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Storage Pool Shortfall Percentage to estimate how critical the shortage of space is for a pool.

This value is only available for pools with thin-provisioned volumes.

Storage Pool Shortfall Percentage

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated.

The higher the percentage, the more critical the shortfall of pool space.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{virtual capacity} - \text{allocated space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage Pool Unallocated Volume Space (GiB)

The amount of the Total Volume Capacity in the pool that is not allocated.

IBM Spectrum Control uses the following formula to determine this value:

$$\text{total volume capacity} - \text{allocated space}$$

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool.

Storage Pool Unused Volume Space (GiB)

The amount of space that is allocated to the volumes in a pool and is not yet used.

IBM Spectrum Control uses the following formula to determine this value:

$$\text{allocated space} - \text{used space}$$

Storage Pool Unallocatable Volume Space (GiB)

The amount of space by which the Total Volume Capacity exceeds the physical capacity of a pool.

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool. If you create volumes with more virtual capacity than can be physically allocated in the pool, you can over commit storage in the pool.

This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

Storage Pool Used Space Percentage

The percentage of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

Storage Pool Compressed Virtual Capacity (GiB)

The total virtual capacity of all the volumes that are compressed in a pool.

Storage Pool Compressed Capacity (GiB)

The amount of storage space that is used by compressed volumes in a pool. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 20.

Storage Pool Uncompressed Capacity (GiB)

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Storage Pool Maximum I/O Capability

The projected maximum number of I/O operations per second for a pool. This value is calculated based on the value in the Storage Pool Back-End Storage Disks property, and on the values in following properties:

- Storage Pool Back-End Storage System Type
- Storage Pool Back-End Storage RAID Level
- Storage Pool Back-End Storage Disk Type

This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage Pool Back-End Storage Disks

The number of physical disks that contribute to the volumes on the back-end storage system. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Component properties

You can create performance reports that include the following information:

Storage Pool Is Encryptable

Shows whether the storage pool can be encrypted. If this value is **Yes**, the storage pool can be encrypted.

Storage Pool Is Encrypted

Shows whether the resource is encrypted. If this value is **Yes**, the resource is encrypted.

Storage Pool Is Solid State

Shows whether there are solid-state drives in the pool. This property can contain the following values:

Mixed

The pool contains both hard disk drives and solid-state drives.

Non solid state

The pool contains no solid-state drives.

Solid state

The pool contains a least one solid-state drive.

Storage Pool Is Thin Provisioned

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is **Yes**, the resource is thin-provisioned.

Storage Pool RAID Level

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Pool Warning Level

The percentage of used capacity of the storage pool at which a warning is generated.

Storage Pool Is Compression Active

Shows whether the compression feature is enabled on the storage pool. If this value is **Yes**, the compression feature is enabled.

Storage Pool Custom Tag 1, 2, and 3

User-defined text that is associated with a storage pool. You can add or edit the custom tags for a storage pool on the Properties notebook of the pool.

Storage Pool Back-End Storage System Type

The type of storage system that provides storage space to a pool. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage Pool Back-End Storage RAID Level

The RAID level of the volumes on the back-end storage system that provide storage space to a pool. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage Pool Back-End Storage Disk Type

The class and speed of the physical disks that contribute to the volumes on the back-end storage system. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage Pool Status

The condition of the resource, for example normal, warning, or error.

Storage Pool Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

- [Performance metrics for pools on SAN Volume Controller and Storwize systems](#)

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for volumes and back-end arrays for pools on these systems in performance reports.

- [Performance metrics for pools on DS8000 storage systems](#)

You can include performance metrics for volumes, caches, back-end arrays, and other data for pools that are on DS8000® storage systems in performance reports.

- [Performance metrics for pools on XIV systems and IBM Spectrum Accelerate](#)

You can include performance data for volumes, caches, and other data for pools on XIV systems and IBM Spectrum Accelerate in performance reports.

Performance metrics for pools on SAN Volume Controller and Storwize systems

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for volumes and back-end arrays for pools on these systems in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Maximum Read I/O Rate (ops/s)

The maximum read I/O rate of the volume in operations per second. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Maximum Back-End Read I/O Rate (ops/s)

The maximum number of read operations per second that are issued to the back-end storage resources. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Maximum Back-End Write I/O Rate (ops/s)

The maximum number of write operations per second that are issued to the back-end storage resources. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)

The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Peak Back-End Read Response Time (ms)

The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)

The longest time for a back-end storage resource to respond to a write operation by a node.

Maximum Back-End Read Response Time (ms/op)

The maximum number of milliseconds for the back-end storage resources to respond to a read operation. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Maximum Back-End Write Response Time (ms/op)

The maximum number of milliseconds for the back-end storage resources to respond to a write operation. This metric is available only for hourly data and daily data for resources that are running IBM Spectrum Virtualize.

Back-End Read Queue Time (ms/op)

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)
The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Performance metrics for pools on DS8000 storage systems

You can include performance metrics for volumes, caches, back-end arrays, and other data for pools that are on DS8000® storage systems in performance reports.

Volume data

You can create performance reports that include the following information:

Normal Read I/O Rate (ops/s)
The average number of nonsequential read operations per second.

Sequential Read I/O Rate (ops/s)
The average number of sequential read operations per second.

Overall Read I/O Rate (ops/s)
The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Normal Write I/O Rate (ops/s)
The average number of nonsequential write operations per second.

Sequential Write I/O Rate (ops/s)
The average number of sequential write operations per second.

Overall Write I/O Rate (ops/s)
The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Normal I/O Rate (ops/s)
The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Sequential I/O Rate (ops/s)
The average number of sequential I/O operations per second. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)
The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Record Mode Read I/O Rate (ops/s)
The average number of I/O operations per second for record-mode read operations.
For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

Read Data Rate (MiB/s)
The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)
The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)
The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)
The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)
The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)
The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)
The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)
The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.
This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Normal Read Cache Hit Percentage
The percentage of nonsequential read operations that find data in the cache.
You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Sequential Read Cache Hit Percentage
The percentage of sequential read operations that find data in the cache.
You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage
The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.
You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Normal Write Cache Hit Percentage
The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage
The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage
The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Normal Cache Hit Percentage
The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.
You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage
The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.
You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage
The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.
You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)
The average number of sectors or tracks per second that are transferred per second from the disks to the cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Cache to Disk Transfer Rate (ops/s)
The average number of sectors or tracks per second that are transferred from the cache to the disks. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Delay Percentage
The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.
This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Record Mode Read Cache Hit Percentage
The percentage of cache hits for record-mode read operations.
For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)
The average number of read operations per second that are issued by the High Performance FICON® feature of the storage system.

HPF Write I/O Rate (ops/s)
The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

PPRC Transfer Rate (ops/s)

The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Performance metrics for pools on XIV systems and IBM Spectrum Accelerate

You can include performance data for volumes, caches, and other data for pools on XIV® systems and IBM Spectrum Accelerate in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Read Data Cache Hit Percentage

The percentage of all read data that is read from the cache.

Write Data Cache Hit Percentage

The percentage of all write data that is written to cache slots that are marked as modified.

Overall Data Cache Hit Percentage

The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

Read Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a read-cache hit operation.

Write Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a write-cache hit operation.

Overall Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

Read Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a read-cache miss operation.

Write Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a write-cache miss operation.

Overall Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

Other data

You can create performance reports that include the following information:

Small Transfers I/O Percentage

The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers I/O Percentage

The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers I/O Percentage

The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers I/O Percentage

The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

Data for managed disks on storage systems in performance reports

You can include general information, capacity data, properties, and other information about managed disks on storage systems in performance reports.

Information about managed disks on storage systems

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Storage System MDisk Name

The name that was assigned to the managed disk on a storage system when it was added to the system.

Capacity and usage data

You can create performance reports that include the following information:

Storage System MDisk Capacity (GiB)

The amount of storage space on the managed disk.

Storage MDisk Available Space (GiB)

The amount of unused storage space on the managed disk on a storage system.

Storage System MDisk Strip Size (KB)

The RAID strip size on a managed disk on a storage system.

Component properties

You can create performance reports that include the following information:

Storage MDisk RAID Level

The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

Storage MDisk Spare Goal

The number of spare drives that are required to maintain redundancy. Use spare drives to protect the system against drive failures in the array on a managed disk on a storage virtualizer.

Storage System MDisk Type

The type of managed disk on a storage system. For example, the disk on a storage system can be a local managed disk.

Storage System MDisk Mode

The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

Array

Extents are provided from local disks.

Managed

Extents are provided from a back-end storage volume.

Unmanaged

The managed disk is not used in the system.

Storage System MDisk Is Balanced

Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is **Yes**, the LUNs are balanced.

Storage System MDisk Fast Write State

Shows whether the cache for a volume on a disk that is managed by a storage system is empty, contains data, or is corrupted.

Storage System MDisk Write Verify

Shows whether all write operations on a managed disk on a storage system are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is **Yes**, all write operations are verified by a follow-up read operation.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage MDisk Status

The condition of the resource, for example normal, warning, or error.

Storage MDisk Native Status

Shows the level of access of nodes in the system to a managed disk or volume on a storage system. The level of access can be online, offline, degraded, or excluded.

- [Performance metrics for managed disks on SAN Volume Controller and Storwize systems](#)

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for back-end arrays on managed disks on these systems in performance reports.

Performance metrics for managed disks on SAN Volume Controller and Storwize systems

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for back-end arrays on managed disks on these systems in performance reports.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)

The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Peak Back-End Read Response Time (ms)

The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)

The longest time for a back-end storage resource to respond to a write operation by a node.

Back-End Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Queue Time (ms/op)

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Data for storage system nodes in performance reports

You can include general information, capacity data, properties, and other information about storage system nodes in performance reports.

Information about storage system nodes

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Node Name

The name that was assigned to the storage system node when it was added to the system.

Storage I/O Group Name

The name that was assigned to the I/O group when it was added to the system.

Component properties

You can create performance reports that include the following information:

Storage Node IP Address

The IP address of the resource.

Storage Node WWN

The worldwide name of the node.

Storage Node Configuration

The number of the node for the resource. This value can be either 0 or 1.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage Node Status

Shows whether a storage system node is online.

- [Performance metrics for nodes on SAN Volume Controller and Storwize systems](#)

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports, front ends, caches, back-end arrays, and other data for nodes on these systems in performance reports.

Performance metrics for nodes on SAN Volume Controller and Storwize systems

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports, front ends, caches, back-end arrays, and other data for nodes on these systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)

The average number of I/O operations per second for operations in which data is sent from a port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive I/O Rate (ops/s)

The average number of I/O operations per second for operations in which the port receives data. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

Port to Host Send I/O Rate (ops/s)

The average number of IOs per second that are sent by the storage system to the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Port to Host Receive I/O Rate (ops/s)
The average number of IOs per second that are received by the storage system from the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Total Port to Host I/O Rate (ops/s)
The average number of IOs per second that are transmitted between the storage system and the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Port to Disk Send I/O Rate (ops/s)
The average number of IOs per second that are sent from the storage system to the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.

Port to Disk Receive I/O Rate (ops/s)
The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)
The average number of IOs per second that are transmitted between the storage system and the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.

Port to Local Node Send I/O Rate (ops/s)
The average number of IOs per second that are sent to other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Port to Local Node Receive I/O Rate (ops/s)
The average number of IOs per second that are received from other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Total Port to Local Node I/O Rate (ops/s)
The average number of IOs per second that are transmitted between the resource and other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Port to Remote Node Send I/O Rate (ops/s)
The average number of IOs per second that are sent to nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Port to Remote Node Receive I/O Rate (ops/s)
The average number of IOs per second that are received from nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Total Port to Remote Node I/O Rate (ops/s)
The average number of IOs per second that are transmitted between the resource and nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Port Send Data Rate (MiB/s)
The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send Data Rate (MiB/s)
The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)
The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)
The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Sph_port_send_bandwidth_percentagend Data Rate (MiB/s)
The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)
The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)
The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Send Data Rate (MiB/s)
The average rate at which data is sent to other nodes that are in the local cluster. The rate is measured in MiB per second.

Port to Local Node Receive Data Rate (MiB/s)
The average rate at which data is received from other nodes that are in the local cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and other nodes that are in the local cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)
The average rate at which data is sent to nodes that are in the remote cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)
The average rate at which data is received from nodes that are in the remote cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and nodes that are in the remote cluster. The rate is measured in MiB per second.

Port to Local Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to another node that is in the local cluster. This value represents the external response time of the transfers.

Port to Local Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from another node that is in the local cluster. This value represents the external response time of the transfers.

Overall Port to Local Node Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation with another node that is in the local cluster. This value represents the external response time of the transfers.

Port to Remote Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to a node that is in the remote cluster. This value represents the external response time of the transfers.

Port to Remote Node Receive Response Time (ms/op)

The average number of milliseconds to complete a receive operation from a node that is in the remote cluster. This value represents the external response time of the transfers.

Overall Port to Remote Node Response Time (ms/op)

The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote cluster. This value represents the external response time of the transfers.

Port to Local Node Send Queue Time (ms/op)

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local cluster.

Port to Local Node Receive Queue Time (ms/op)

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local cluster.

Overall Port to Local Node Queue Time (ms/op)

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local cluster.

Port to Remote Node Send Queue Time (ms/op)

The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote cluster.

Port to Remote Node Receive Queue Time (ms/op)

The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote cluster.

Overall Port to Remote Node Queue Time (ms/op)

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote cluster.

Front end data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)

The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Overflow I/O Rate (ops/s)

The average number of tracks per second that are written but are delayed because there is not enough space in the write cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Flush Through I/O Rate (ops/s)

The average number of tracks per second that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Write Through I/O Rate (ops/s)

The average number of tracks per second that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable. You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Write Cache Delay Percentage

The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Read Ahead Percentage of Cache Hits

The percentage of all read cache hits that occur on pre-staged data. This value applies only to the volume copy cache if the resource is running IBM Spectrum Virtualize 7.3 or later.

Dirty Write Percentage of Cache Hits

The percentage of all cache write hits that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Overflow Percentage

The percentage of write operations that are delayed because there is not enough space in the write cache.

Write Cache Flush Through Percentage

The percentage of tracks that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Write Through Percentage

The percentage of tracks that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Disk to Cache Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred per second from the disks to the cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Cache to Disk Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred from the cache to the disks. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Back-end data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)

The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Peak Back-End Read Response Time (ms)

The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)

The longest time for a back-end storage resource to respond to a write operation by a node.

Back-End Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Queue Time (ms/op)

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)

The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write Percentage

The percentage of overlapping write operations that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.

Applies to resources that are running IBM Spectrum Virtualize.

Global Mirror Overlapping Write I/O Rate (op/s)

The average number of overlapping write operations per second that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.

This value applies to resources that are running IBM Spectrum Virtualize.

Global Mirror Secondary Write Lag (ms/op)

The average number of additional milliseconds that it takes to service each secondary write operation for Global Mirror. This value does not include the time to service the primary write operations.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Average number of additional milliseconds it took to service each secondary write operation for Global Mirror, beyond the time needed to service the primary writes

Peak Read Response Time (ms)

The worst response time measured for a read operation in the sample interval.

Peak Write Response Time (ms)

The worst response time measured for a write operation in the sample interval.

Link Failure Rate (count/s)

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

Loss of Sync Rate (count/s)

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

CRC Error Rate (count/s)

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.

Primitive Sequence Protocol Error Rate (count/s)

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)

The average number of bit errors per second that are detected.

Processor Utilization Percentage

The average percentage of time that the processors on nodes are busy doing system I/O tasks. This value applies only to resources that are running IBM Spectrum Virtualize.

Overall Host Attributed Response Time Percentage

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Zero Buffer-to-Buffer Credit Timer (microseconds)

The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that metadata was collected.

If this metric is not available, use the Port Send Delay Time metric instead.

Zero Buffer-to-Buffer Credit Percentage

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that metadata was collected. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

Data for storage system I/O groups in performance reports

You can include general information, capacity data, properties, and other information about storage system I/O groups in performance reports.

Information about storage system I/O groups

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage I/O Group Name

The name that was assigned to the I/O group when it was added to the system.

Component properties

You can create performance reports that include the following information:

Storage I/O Group Mirroring Total Memory (MiB)

The total amount of memory that is allocated on the nodes in an I/O group to provide volume mirroring functions.

Storage I/O Group Mirroring Free Memory (MiB)

The amount of available memory on the nodes in an I/O group to provide volume mirroring functions.

Storage I/O Group FlashCopy Total Memory (MiB)

The total amount of memory that is allocated on the nodes in an I/O group to provide FlashCopy® mirroring functions.

Storage I/O Group FlashCopy Free Memory (MiB)

The amount of available memory on the nodes in an I/O group to provide FlashCopy mirroring functions.

Storage I/O Group Remote Copy Total Memory (MiB)

The total amount of memory that is allocated on the nodes in an I/O group to provide mirroring functions by using a remote copy.

Storage I/O Group Remote Copy Free Memory (MiB)

The amount of available memory on the nodes in an I/O group to provide mirroring functions by using a remote copy.

Storage I/O Group RAID Total Memory (MiB)

The total amount of memory that is allocated on the nodes in an I/O group to provide RAID mirroring functions.

Storage I/O Group RAID Free Memory (MiB)

The amount of available memory on the nodes in an I/O group to provide RAID mirroring functions.

Status information

You can create performance reports that include the following information:

Storage I/O Group Maintenance

Shows whether the I/O group is in maintenance mode. If this value is **Yes**, the I/O group is in maintenance mode.

Storage I/O Group Is Compression Active

Shows whether the compression feature is enabled on the volumes in the storage I/O group. If this value is **Yes**, the compression feature is enabled.

Storage I/O Group Is Compression Supported

Shows whether the compression feature is available for volumes in the storage I/O group. If this value is **Yes**, the compression feature is available.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

- **Performance metrics for I/O groups on SAN Volume Controller and Storwize systems**

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports, volumes, caches, back-end arrays, and other data for I/O groups on these systems in performance reports.

Performance metrics for I/O groups on SAN Volume Controller and Storwize systems

A storage environment can include SAN Volume Controller systems, Storwize® V7000 systems, and a Storwize V7000 Unified systems. You can include performance metrics for ports, volumes, caches, back-end arrays, and other data for I/O groups on these systems in performance reports.

Port data

You can create performance reports that include the following information:

Port Send I/O Rate (ops/s)

The average number of I/O operations per second for operations in which data is sent from a port. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive I/O Rate (ops/s)

The average number of I/O operations per second for operations in which the port receives data. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port I/O Rate (ops/s)

The average number of send operations and receive operations per second.

Port to Host Send I/O Rate (ops/s)

The average number of IOs per second that are sent by the storage system to the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Port to Host Receive I/O Rate (ops/s)

The average number of IOs per second that are received by the storage system from the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Total Port to Host I/O Rate (ops/s)

The average number of IOs per second that are transmitted between the storage system and the hosts that are accessing its storage. Use this metric to help measure host workload against the storage system.

Port to Disk Send I/O Rate (ops/s)

The average number of IOs per second that are sent from the storage system to the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.

Port to Disk Receive I/O Rate (ops/s)

The average number of exchanges per second that are received from back-end storage resources.

Total Port to Disk I/O Rate (ops/s)

The average number of IOs per second that are transmitted between the storage system and the back-end storage it is virtualizing. Use this metric to help measure the rate of data that is sent to back-end storage.

Port to Local Node Send I/O Rate (ops/s)

The average number of IOs per second that are sent to other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Port to Local Node Receive I/O Rate (ops/s)

The average number of IOs per second that are received from other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Total Port to Local Node I/O Rate (ops/s)
The average number of I/Os per second that are transmitted between the resource and other nodes within the local cluster. Use this metric to understand the rate of inter-cluster communication.

Port to Remote Node Send I/O Rate (ops/s)
The average number of I/Os per second that are sent to nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Port to Remote Node Receive I/O Rate (ops/s)
The average number of I/Os per second that are received from nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Total Port to Remote Node I/O Rate (ops/s)
The average number of I/Os per second that are transmitted between the resource and nodes that are in a remote cluster. Use this metric to understand the amount of remote replication workload.

Port Send Data Rate (MiB/s)
The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)
The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)
The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port to Host Send Data Rate (MiB/s)
The average rate at which data is sent to host computers. The rate is measured in MiB per second.

Port to Host Receive Data Rate (MiB/s)
The average rate at which data is received from host computers. The rate is measured in MiB per second.

Total Port to Host Data Rate (MiB/s)
The average rate at which data is transmitted between host computers and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Disk Sph_port_send_bandwidth_percentageend Data Rate (MiB/s)
The average rate at which data is sent to back-end storage resources. The rate is measured in MiB per second.

Port to Disk Receive Data Rate (MiB/s)
The average rate at which data is received from back-end storage resources. The rate is measured in MiB per second.

Total Port to Disk Data Rate (MiB/s)
The average rate at which data is transmitted between back-end storage resources and the component. The rate is measured in MiB per second and includes both send and receive operations.

Port to Local Node Send Data Rate (MiB/s)
The average rate at which data is sent to other nodes that are in the local cluster. The rate is measured in MiB per second.

Port to Local Node Receive Data Rate (MiB/s)
The average rate at which data is received from other nodes that are in the local cluster. The rate is measured in MiB per second.

Total Port to Local Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and other nodes that are in the local cluster. The rate is measured in MiB per second.

Port to Remote Node Send Data Rate (MiB/s)
The average rate at which data is sent to nodes that are in the remote cluster. The rate is measured in MiB per second.

Port to Remote Node Receive Data Rate (MiB/s)
The average rate at which data is received from nodes that are in the remote cluster. The rate is measured in MiB per second.

Total Port to Remote Node Data Rate (MiB/s)
The average rate at which data is transmitted between the component and nodes that are in the remote cluster. The rate is measured in MiB per second.

Port to Local Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to another node that is in the local cluster. This value represents the external response time of the transfers.

Port to Local Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from another node that is in the local cluster. This value represents the external response time of the transfers.

Overall Port to Local Node Response Time (ms/op)
The average number of milliseconds to complete a send or receive operation with another node that is in the local cluster. This value represents the external response time of the transfers.

Port to Remote Node Send Response Time (ms/op)
The average number of milliseconds to complete a send operation to a node that is in the remote cluster. This value represents the external response time of the transfers.

Port to Remote Node Receive Response Time (ms/op)
The average number of milliseconds to complete a receive operation from a node that is in the remote cluster. This value represents the external response time of the transfers.

Overall Port to Remote Node Response Time (ms/op)
The average number of milliseconds to complete a send operation to, or a receive operation from a node in the remote cluster. This value represents the external response time of the transfers.

Port to Local Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to other nodes that are in the local cluster.

Port to Local Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from other nodes that are in the local cluster.

Overall Port to Local Node Queue Time (ms/op)
The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and other nodes that are in the local cluster.

Port to Remote Node Send Queue Time (ms/op)
The average time in milliseconds that a send operation spends in the queue before the operation is processed. This value represents the queue time for send operations that are issued to a node that is in the remote cluster.

Port to Remote Node Receive Queue Time (ms/op)
The average time in milliseconds that a receive operation spends in the queue before the operation is processed. This value represents the queue time for receive operations that are issued from a node that is in the remote cluster.

Overall Port to Remote Node Queue Time (ms/op)

The average number of milliseconds that a send or receive operation spends in the queue before the operation is processed. This value is for send and receive operations that are issued between the component and a node that is in the remote cluster.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)

The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)

The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)

The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)

The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)

The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)

The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)

The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Peak Back-End Read Response Time (ms)

The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)

The longest time for a back-end storage resource to respond to a write operation by a node.

Back-End Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)

The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Queue Time (ms/op)

The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)

The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)

The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Write Queue Time (ms)

The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Write Cache Delay Percentage

The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Flush Through I/O Rate (ops/s)

The average number of tracks per second that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Write Through I/O Rate (ops/s)

The average number of tracks per second that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Overflow I/O Rate (ops/s)

The average number of tracks per second that are written but are delayed because there is not enough space in the write cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Delay I/O Rate (ops/s)

The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Disk to Cache Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred per second from the disks to the cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer.

Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Cache to Disk Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred from the cache to the disks. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Read Ahead Percentage of Cache Hits

The percentage of all read cache hits that occur on pre-staged data. This value applies only to the volume copy cache if the resource is running IBM Spectrum Virtualize 7.3 or later.

Dirty Write Percentage of Cache Hits

The percentage of all cache write hits that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Flush Through Percentage

The percentage of tracks that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Write Through Percentage

The percentage of tracks that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Overflow Percentage

The percentage of write operations that are delayed because there is not enough space in the write cache.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)

The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write Percentage

The percentage of overlapping write operations that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.

Applies to resources that are running IBM Spectrum Virtualize.

Global Mirror Overlapping Write I/O Rate (op/s)

The average number of overlapping write operations per second that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.

This value applies to resources that are running IBM Spectrum Virtualize.

Peak Read Response Time (ms)

The worst response time measured for a read operation in the sample interval.

Peak Write Response Time (ms)

The worst response time measured for a write operation in the sample interval.

Global Mirror Secondary Write Lag (ms/op)

The average number of additional milliseconds that it takes to service each secondary write operation for Global Mirror. This value does not include the time to service the primary write operations.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Average number of additional milliseconds it took to service each secondary write operation for Global Mirror, beyond the time needed to service the primary writes

Link Failure Rate (count/s)

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

Loss of Sync Rate (count/s)

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

CRC Error Rate (count/s)

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.

Processor Utilization Percentage

The average percentage of time that the processors on nodes are busy doing system I/O tasks. This value applies only to resources that are running IBM Spectrum Virtualize.

Overall Host Attributed Response Time Percentage

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Nonpreferred Node Usage Percentage

The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.

Primitive Sequence Protocol Error Rate (count/s)

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)

The average number of bit errors per second that are detected.

Zero Buffer-to-Buffer Credit Timer (microseconds)

The number of microseconds that the port is not able to send frames between ports because there is insufficient buffer-to-buffer credit.

In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports. Buffer-to-buffer credit is measured from the last time that metadata was collected.

If this metric is not available, use the Port Send Delay Time metric instead.

Zero Buffer-to-Buffer Credit Percentage

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that metadata was collected. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

Data for local disks on storage systems in performance reports

You can include general information, capacity data, properties, and other information about local disks that are on storage systems in performance reports.

Information about local disks on storage systems

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage MDisk Name

The name that was assigned to the managed disk when it was added to the system.

Storage Local Disk Name

The name that was assigned to the local disk on the storage system when the local disk was created on the system.

Capacity and usage data

You can create performance reports that include the following information:

Storage Local Disk Capacity (GiB)

The amount of storage space on the local disk.

Component properties

You can create performance reports that include the following information:

Storage Local Disk Is Encryptable

Shows whether the resource can be encrypted. If this value is **Yes**, the resource can be encrypted.

Storage Local Disk Is Encrypted

Shows whether the resource is encrypted. If this value is **Yes**, the resource is encrypted.

Storage Local Disk Storage Class

The storage technology of the local disk on a storage system. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD).

Storage Local Disk Is Solid State

Shows whether the disk is a solid-state drive.

Storage Local Disk Speed (RPM)
The speed of the local disk on a storage system.

Storage System Configuration
Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Status information

You can create performance reports that include the following information:

Storage Local Disk Status
The condition of the resource, for example normal, warning, or error.

Vendor, model, and device information

You can create performance reports that include the following information:

Storage System Local Disk Firmware Version
The version number of the firmware that is running on the disk.

Storage System Local Disk Model
The model name or model number of the resource.

Storage System Local Disk Serial Number
The serial number of the resource.

Storage System Local Disk Vendor
The vendor who supplied the resource.

- [Performance metrics for local disks on storage systems](#)
You can include performance data for back-end arrays on local disks that are on storage systems in performance reports.

Performance metrics for local disks on storage systems

You can include performance data for back-end arrays on local disks that are on storage systems in performance reports.

Back-end array data

You can create performance reports that include the following information:

Back-End Read I/O Rate (ops/s)
The average number of read operations per second that are issued to the back-end storage resources.

Back-End Write I/O Rate (ops/s)
The average number of write operations per second that are issued to the back-end storage resources.

Total Back-End I/O Rate (ops/s)
The average number of I/O operations per second that are transmitted between the back-end storage resources and the component. This value includes both read and write operations.

Back-End Read Data Rate (MiB/s)
The average number of MiB per second that are read from the back-end storage resources.

Back-End Write Data Rate (MiB/s)
The average number of MiB per second that are written to the back-end storage resources.

Total Back-End Data Rate (MiB/s)
The average rate at which data is transmitted between the back-end storage resources and the component. The rate is measured in MiB per second and includes both read and write operations.

Back-End Read Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read operation.

Back-End Write Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a write operation.

Overall Back-End Response Time (ms/op)
The average number of milliseconds for the back-end storage resources to respond to a read or a write operation.

Peak Back-End Read Response Time (ms)
The longest time for a back-end storage resource to respond to a read operation.

Peak Back-End Write Response Time (ms)
The longest time for a back-end storage resource to respond to a write operation by a node.

Back-End Read Transfer Size (KiB/op)
The average number of KiB that are transferred per read operation from the back-end storage resources.

Back-End Write Transfer Size (KiB/op)
The average number of KiB that are transferred per write operation to the back-end storage resources.

Overall Back-End Transfer Size (KiB/op)
The average transfer size, in KiB, of a read or a write operation to the back-end storage resources.

Back-End Read Queue Time (ms/op)
The average number of milliseconds that a read operation spends in the queue before the operation is sent to the back-end storage resources.

Back-End Write Queue Time (ms/op)
The average number of milliseconds that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Overall Back-End Queue Time (ms/op)
The average number of milliseconds that a read or a write operation spends in the queue before the operation is sent to the back-end storage resources.

Peak Back-End Read Queue Time (ms)

The longest time that a read operation spends in the queue before the operation is sent to the back-end storage resources.
Peak Back-End Write Queue Time (ms)
The longest time that a write operation spends in the queue before the operation is sent to the back-end storage resources.

Data for storage host connections in performance reports

You can include general information, capacity data, properties, and other information about storage host connections in performance reports.

Information about storage host connections

You can create performance reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Storage Host Connection Name

The user-defined name that describes the server or storage device that is assigned to a volume.

XIV® Cluster

The name of the cluster, if any, as defined on the XIV. If no cluster is defined for the host connection, this field is blank.

- [Performance metrics for host connections on SAN Volume Controller, Storwize V7000, or Storwize V7000 Unified systems](#)
You can include performance metrics for volumes, caches, and for other data for host connections in performance reports.
- [Performance metrics for host connections on DS8000 storage systems](#)
You can include performance metrics for volumes, caches, and for other data for storage host connections on DS8000® storage systems in performance reports.
- [Performance metrics for storage host connections on an XIV](#)
You can include performance metrics for volumes, caches, and for other data for storage host connections on an XIV in performance reports.

Performance metrics for host connections on SAN Volume Controller, Storwize® V7000, or Storwize V7000 Unified systems

You can include performance metrics for volumes, caches, and for other data for host connections in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Disk to Cache Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred per second from the disks to the cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Cache to Disk Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred from the cache to the disks. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Delay Percentage

The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Flush Through I/O Rate (ops/s)

The average number of tracks per second that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Write Through I/O Rate (ops/s)

The average number of tracks per second that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Overflow I/O Rate (ops/s)

The average number of tracks per second that are written but are delayed because there is not enough space in the write cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Delay I/O Rate (ops/s)

The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Read Ahead Percentage of Cache Hits

The percentage of all read cache hits that occur on pre-staged data. This value applies only to the volume copy cache if the resource is running IBM Spectrum Virtualize 7.3 or later.

Dirty Write Percentage of Cache Hits

The percentage of all cache write hits that occur on data in the cache that is marked as modified. This value represents how effectively write operations are coalesced before the data is written to disk. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Flush Through Percentage

The percentage of tracks that are written to disk in flush-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Write Through Percentage

The percentage of tracks that are written to disk in write-through mode. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3.

Write Cache Overflow Percentage

The percentage of write operations that are delayed because there is not enough space in the write cache.

Other data

You can create performance reports that include the following information:

Global Mirror Write I/O Rate (op/s)

The average number of write operations per second that are issued to the Global Mirror secondary site.

Global Mirror Overlapping Write I/O Rate (op/s)

The average number of overlapping write operations per second that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.

This value applies to resources that are running IBM Spectrum Virtualize.

Peak Read Response Time (ms)

The worst response time measured for a read operation in the sample interval.

Peak Write Response Time (ms)

The worst response time measured for a write operation in the sample interval.

Global Mirror Secondary Write Lag (ms/op)

The average number of additional milliseconds that it takes to service each secondary write operation for Global Mirror. This value does not include the time to service the primary write operations.

You monitor the value of Global Mirror Secondary Write Lag to identify delays that occurred during the process of writing data to the secondary site.

Average number of additional milliseconds it took to service each secondary write operation for Global Mirror, beyond the time needed to service the primary writes

Overall Host Attributed Response Time Percentage

The percentage of the average response time that can be attributed to delays from host systems. This value includes both read response times and write response times, and can help you diagnose slow hosts and fabrics that are not working efficiently.

For read response time, the value is based on the time that it takes for hosts to respond to transfer-ready notifications from the nodes. For write response time, the value is based on the time that it takes for hosts to send the write data after the node responds to a transfer-ready notification.

Global Mirror Overlapping Write Percentage

The percentage of overlapping write operations that are issued by the Global Mirror primary site. Some overlapping writes are processed in parallel and are excluded from this value.

Applies to resources that are running IBM Spectrum Virtualize.

Nonpreferred Node Usage Percentage

The overall percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume.

Performance metrics for host connections on DS8000 storage systems

You can include performance metrics for volumes, caches, and for other data for storage host connections on DS8000® storage systems in performance reports.

Volume data

You can create performance reports that include the following information:

Normal Read I/O Rate (ops/s)

The average number of nonsequential read operations per second.

Sequential Read I/O Rate (ops/s)

The average number of sequential read operations per second.

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Normal Write I/O Rate (ops/s)

The average number of nonsequential write operations per second.

Sequential Write I/O Rate (ops/s)

The average number of sequential write operations per second.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Normal I/O Rate (ops/s)

The average number of nonsequential I/O operations per second. This value includes both read and write operations.

Total Sequential I/O Rate (ops/s)

The average number of sequential I/O operations per second. This value includes both read and write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Record Mode Read I/O Rate (ops/s)

The average number of I/O operations per second for record-mode read operations.

For record-mode read operations, only the requested data is managed in the cache rather than a full track of data.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Write Cache Delay I/O Rate (ops/s)

The average number of I/O operations per second that are delayed because of space constraints in the write cache, or because of other conditions.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Normal Read Cache Hit Percentage

The percentage of nonsequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Sequential Read Cache Hit Percentage

The percentage of sequential read operations that find data in the cache.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Normal Write Cache Hit Percentage

The percentage of nonsequential write operations that are handled in the cache.

Sequential Write Cache Hit Percentage

The percentage of sequential write operations that are handled in the cache.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Normal Cache Hit Percentage

The percentage of nonsequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Sequential Cache Hit Percentage

The percentage of sequential I/O operations that are handled in the cache. This value includes both read and write operations.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Record Mode Read Cache Hit Percentage

The percentage of cache hits for record-mode read operations.

For record-mode read operations, only the requested data, rather than a full track of data, is managed in the cache.

Disk to Cache Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred per second from the disks to the cache. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer.

Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Cache to Disk Transfer Rate (ops/s)

The average number of sectors or tracks per second that are transferred from the cache to the disks. This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Write Cache Delay Percentage

The percentage of I/O operations that are delayed because of space constraints in the write cache, or because of other conditions. The value is a percentage of all operations.

This value applies only to resources that are running a version of IBM Spectrum Virtualize earlier than 7.3. The metric is an approximation because actual transfer rates are different for each cache layer. Therefore, you cannot directly compare this metric from resources that are 7.3 or later with earlier versions.

Other data

You can create performance reports that include the following information:

HPF Read I/O Rate (ops/s)

The average number of read operations per second that are issued by the High Performance FICON® feature of the storage system.

HPF Write I/O Rate (ops/s)

The average number of write operations per second that are issued by the High Performance FICON feature of the storage system.

Total HPF I/O Rate (ops/s)

The average number of I/O operations per second that are issued by the High Performance FICON feature of the storage system. This value includes both read and write operations.

HPF I/O Percentage

The percentage of all I/O operations that are issued by the High Performance FICON feature of the storage system.

PPRC Transfer Rate (ops/s)

The average number of tracks per second that are transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair. This value shows the activity for the source of the PPRC relationship, but shows no activity for the target.

Performance metrics for storage host connections on an XIV

You can include performance metrics for volumes, caches, and for other data for storage host connections on an XIV® in performance reports.

Volume data

You can create performance reports that include the following information:

Overall Read I/O Rate (ops/s)

The average number of read operations per second. This value includes both sequential and nonsequential read operations.

Overall Write I/O Rate (ops/s)

The average number of write operations per second. This value includes both sequential and nonsequential write operations.

Total Overall I/O Rate (ops/s)

The average number of read operations and write operations per second. This value includes both sequential and nonsequential operations.

Read Data Rate (MiB/s)

The average number of MiBs per second that are transferred for read operations.

Write Data Rate (MiB/s)

The average number of MiBs per second that are transferred for write operations.

Total Data Rate (MiB/s)

The average number of MiB per second that are transferred for read operations and write operations.

Read Response Time (ms/op)

The average number of milliseconds to complete a read operation.

Write Response Time (ms/op)

The average number of milliseconds to complete a write operation.

Overall Response Time (ms/op)

The average number of milliseconds to complete an I/O operation. This value includes both read and write operations.

Read Transfer Size (KiB/op)

The average number of KiB that are transferred per read operation.

Write Transfer Size (KiB/op)

The average number of KiB that are transferred per write operation.

Overall Transfer Size (KiB/op)

The average number of KiB that are transferred per I/O operation. This value includes both read and write operations.

Cache data

You can create performance reports that include the following information:

Overall Read Cache Hit Percentage

The percentage of all read operations that find data in the cache. This value includes both sequential and random read operations, and read operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because in the event of a cache miss, the data must be read from the back-end storage resources.

Overall Write Cache Hit Percentage

The percentage of all write operations that are handled in the cache. This value includes both sequential and random write operations, and write operations in the volume cache and volume copy cache where applicable.

Total Overall Cache Hit Percentage

The percentage of all read operations and write operations that are handled in the cache. This value includes both sequential and random read and write operations, and read and write operations in the volume cache and volume copy cache where applicable.

You can use this value to understand throughput or response times. Low cache-hit percentages can increase response times because a cache miss requires access to the back-end storage resources.

Read Data Cache Hit Percentage

The percentage of all read data that is read from the cache.

Write Data Cache Hit Percentage

The percentage of all write data that is written to cache slots that are marked as modified.

Overall Data Cache Hit Percentage

The percentage of all data that is handled in the cache. This value includes read data that is read from the cache and write data that is written to cache slots that are marked as modified.

Read Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a read-cache hit operation.

Write Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a write-cache hit operation.

Overall Cache Hit Response Time (ms/op)

The average number of milliseconds to complete a cache hit operation. This value includes the times for both read-cache hit and write-cache hit operations.

Read Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a read-cache miss operation.

Write Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a write-cache miss operation.

Overall Cache Miss Response Time (ms/op)

The average number of milliseconds to complete a cache miss operation. This value includes the times for both read-cache miss and write-cache miss operations.

Other data

You can create performance reports that include the following information:

Small Transfers I/O Percentage

The percentage of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers I/O Percentage

The percentage of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers I/O Percentage

The percentage of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers I/O Percentage

The percentage of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Data Percentage

The percentage of data that is transferred as a result of I/O operations with a data transfer size that is greater than 512 KiB.

Small Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a small data transfer size. A small data transfer has a size that is less than or equal to 8 KiB.

Medium Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a medium data transfer size. A medium data transfer has a size that is greater than 8 KiB and less than or equal to 64 KiB.

Large Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation that has a large data transfer size. A large data transfer has a size that is greater than 64 KiB and less than or equal to 512 KiB.

Very Large Transfers Response Time (ms/op)

The average number of milliseconds to complete an I/O operation with a data transfer size that is greater than 512 KiB.

Data for switches in performance reports

You can include properties and other information about switches in performance reports.

Information about switches

You can create performance reports that include the following information:

Switch Logical Name

The logical name of a switch or switch port.

Component properties

You can create performance reports that include the following information:

Switch WWN

The World Wide Name (WWN) of the switch. A WWN is the unique 64-bit identifier for the switch.

Parent Switch WWN

The worldwide name of the parent switch.

Switch Management Telnet Address

The Telnet address that you access to manage the switch.

Switch Management SNMP Address

The Simple Network Management Protocol (SNMP) address that you access to manage the switch.

Switch Management URL Address

The URL that you access to manage the switch.

Switch Domain

The domain ID of a switch. The ID is an 8-bit identifier with a range of 0-255. This column is blank for physical switches that are parents of virtual switches.

Switch Version

The version number of the operating system that is running on the switch.

Switch Location

The physical location of the switch. The location is defined when a switch is added to IBM Spectrum® Control. You can add or edit the location of the switch on the General tab of the properties notebook.

Switch Custom Tag 1, 2, and 3

User-defined text that is associated with the switch. You can add or edit the custom tags for the switch on the General tab of the properties notebook.

Status information

You can create performance reports that include the following information:

Switch Status

The condition of the resource, for example normal, warning, or error.

Switch Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Vendor, model, and device information

You can create performance reports that include the following information:

Switch IP Address

The IP address of the resource.

Switch Serial Number

The serial number of the resource.

Switch Vendor

The vendor who supplied the resource.

Switch Model

The model name or model number of the resource.

- [Performance metrics for switches](#)

You can include performance data for switches, including switch error data, in performance reports.

Performance metrics for switches

You can include performance data for switches, including switch error data, in performance reports.

Switch data

You can create performance reports that include the following information:

Port Send Frame Rate (frames/s)

The average number of frames per second that are sent by the port.

Port Receive Frame Rate (frames/s)

The average number of frames per second that are received by the port.

Total Port Frame Rate (frames/s)

The average number of frames per second that are transferred. This value includes frames that are sent and received by the port.

Port Send Data Rate (MiB/s)

The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)

The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Send Frame Size (KiB/op)

The average size of a frame, in KiB, that is sent through the port.

Port Receive Frame Size (KiB/op)

The average size of a frame, in KiB, that is received by the port.

Overall Port Frame Size (KiB/op)

The average frame transfer size. This value is measured in KiB and includes frames that are sent and frames that are received by the port.

Port Send Bandwidth Percentage

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage

The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Switch error data

You can create performance reports that include the following information:

Error Frame Rate (count/s)

The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

Discarded Frame Rate (count/s)

The average number of frames per second that are discarded because host buffers are unavailable for the port.

Link Failure Rate (count/s)

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

Loss of Sync Rate (count/s)

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled.

Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

CRC Error Rate (count/s)

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.

Encoding Disparity Error Rate (count/s)

The average number of disparity errors per second that are received.

Discarded Class 3 Frame Rate (count/s)

The average number of class 3 frames per second that are discarded.

Primitive Sequence Protocol Error Rate (count/s)

The average number of primitive sequence protocol errors per second that are detected.

This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)

The average number of bit errors per second that are detected.

Link Reset Transmitted Rate (count/s)

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

Link Reset Received Rate (count/s)

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

Short Frame Rate (count/s)

The average number of frames that are received per second that are shorter than 28 octets. This number excludes start-of-frame bytes and end-of-frame bytes.

The 28 octet limit is calculated based on the assumption that a frame has 24 header bytes, and 4 CRC bytes.

Long Frame Rate (count/s)

The average number of frames that are received per second that are longer than 2140 octets. This number excludes start-of-frame bytes and end-of-frame bytes.

The 2140 octet limit is calculated based on the assumption that a frame has 24 header bytes, 4 CRC bytes, and 2112 data bytes.

F-BSY Frame Rate (count/s)

The average number of F-BSY frames per second that are generated.

An F-BSY frame is issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N_port is busy.

F-RJT Frame Rate (count/s)

The average number of F-RJT frames per second that are generated.

An F-RJT frame is issued by the fabric to indicate that delivery of a frame was denied.

Zero Buffer Credit Rate

The average number of Zero Buffer Credit conditions per second that occurred. A Zero Buffer Credit condition occurs when a port is unable to send frames because of a lack of buffer credit since the last node reset. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

Class 3 Send Timeout Frame Rate

The average number of class 3 frames per second that were discarded before transmission because of a timeout condition. The timeout condition occurs while the switch or port waits for buffer credit from the receiving port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

Class 3 Receive Timeout Frame Rate

The average number of class 3 frames per second that were discarded after reception because of a timeout condition. The timeout condition occurs while a transmitting port waits for buffer credit from a port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

RDY Priority Override Rate

The average number of times per second during which the sending of R_RDY or VC_RDY signals was a higher priority than the sending of frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

Port State Change Rate

The average number of times per second that the state of a port changes to offline, online, or faulty. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

Port Congestion Index

The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

Zero Buffer-to-Buffer Credit Percentage

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that metadata was collected. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

Credit Recovery Link Reset Rate

The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

Data for switch ports in performance reports

You can include properties and other information about switch ports in performance reports.

Information about switch ports

You can create performance reports that include the following information:

Switch Logical Name

The logical name of a switch or switch port.

Switch Port WWPN

The worldwide port name of the port on the switch.

Component properties

You can create performance reports that include the following information:

Switch Port Number

The port number on the resource.

Switch Port Type

The type of port on the storage system, storage virtualizer, or switch. For example, the port type can be N_Port, F_Port, or another type of port.

Switch Port Speed (GiB/s)

The speed of a port, which is measured in GiB per second.

Switch Blade Slot Number

The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

Status information

You can create performance reports that include the following information:

Switch Port Status

The condition of the resource, for example normal, warning, or error.

Switch Port Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Switch Port Enabled State

Shows whether a port is enabled, disabled, or is enabled but offline.

Switch Port Operational Status

The current operational state of the switch port. For example, this value can be **Error** or **OK**.

- [Performance metrics for switch ports](#)

You can include performance data for switch ports, including port error data, in performance reports.

Performance metrics for switch ports

You can include performance data for switch ports, including port error data, in performance reports.

Switch port data

You can create performance reports that include the following information:

Port Send Frame Rate (frames/s)

The average number of frames per second that are sent by the port.

Port Receive Frame Rate (frames/s)

The average number of frames per second that are received by the port.

Total Port Frame Rate (frames/s)

The average number of frames per second that are transferred. This value includes frames that are sent and received by the port.

Port Send Data Rate (MiB/s)

The average rate at which data is sent from the port. The rate is measured in MiB per second. A send operation is a read operation that is processed, or a write operation that is initiated by the port.

Port Receive Data Rate (MiB/s)

The average rate at which data is received by the port. The rate is measured in MiB per second. A receive operation is a write operation that is processed, or a read operation that is initiated by the port.

Total Port Data Rate (MiB/s)

The average rate at which data is transferred through the port. The rate is measured in MiB per second and includes both send and receive operations.

Port Peak Send Data Rate (MiB/s)

The fastest rate at which data is sent through this switch port.

Port Peak Receive Data Rate (MiB/s)

The fastest rate at which data is received through this switch port.

Port Send Frame Size (KiB/op)

The average size of a frame, in KiB, that is sent through the port.

Port Receive Frame Size (KiB/op)

The average size of a frame, in KiB, that is received by the port.

Overall Port Frame Size (KiB/op)

The average frame transfer size. This value is measured in KiB and includes frames that are sent and frames that are received by the port.

Port Send Bandwidth Percentage

The percentage of the port bandwidth that is used for send operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Port Receive Bandwidth Percentage

The percentage of the port bandwidth that is used for receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Overall Port Bandwidth Percentage

The percentage of the port bandwidth that is used for send and receive operations. This value is an indicator of port bandwidth usage that is based on the speed of the port.

Data for errors on switch ports

You can create performance reports that include the following information:

Error Frame Rate (count/s)

The average number of error frames per second that are received. An error frame is a frame that violates the Fibre Channel Protocol.

Discarded Frame Rate (count/s)

The average number of frames per second that are discarded because host buffers are unavailable for the port.

Link Failure Rate (count/s)

The average number of miscellaneous fibre channel link errors per second for ports. Link errors might occur when an unexpected Not Operational (NOS) is received or a link state machine failure was detected.

Loss of Sync Rate (count/s)

The average number of times per second that the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled. Synchronization is assumed lost after a timeout interval expires.

Loss of Signal Rate (count/s)

The average number of times per second at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However, in some cases, this error can also occur when the maximum link distance between ports is exceeded, for the type of connecting cable and light source.

CRC Error Rate (count/s)

The average number of frames per second that are received in which a cyclic redundancy check (CRC) error is detected. A CRC error is detected when the CRC in the transmitted frame does not match the CRC computed by the receiver. For Brocade switches, this metric includes only the CRC Errors with a good end-of-frame (EOF) indicator.

Encoding Disparity Error Rate (count/s)

The average number of disparity errors per second that are received.

Discarded Class 3 Frame Rate (count/s)

The average number of class 3 frames per second that are discarded.

Primitive Sequence Protocol Error Rate (count/s)

The average number of primitive sequence protocol errors per second that are detected. This error occurs when there is a link failure for a port.

Invalid Transmission Word Rate (count/s)

The average number of bit errors per second that are detected.

Link Reset Transmitted Rate (count/s)

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR1) state.

Link Reset Received Rate (count/s)

The average number of times per second that the port changes from an active (AC) state to a Link Recovery (LR2) state.

Short Frame Rate (count/s)

The average number of frames that are received per second that are shorter than 28 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 28 octet limit is calculated based on the assumption that a frame has 24 header bytes, and 4 CRC bytes.

Long Frame Rate (count/s)

The average number of frames that are received per second that are longer than 2140 octets. This number excludes start-of-frame bytes and end-of-frame bytes. The 2140 octet limit is calculated based on the assumption that a frame has 24 header bytes, 4 CRC bytes, and 2112 data bytes.

F-BSY Frame Rate (count/s)

The average number of F-BSY frames per second that are generated.

An F-BSY frame is issued by the fabric to indicate that a frame cannot be delivered because the fabric or destination N_port is busy.

F-RJT Frame Rate (count/s)

The average number of F-RJT frames per second that are generated.

An F-RJT frame is issued by the fabric to indicate that delivery of a frame was denied.

Zero Buffer Credit Rate

The average number of Zero Buffer Credit conditions per second that occurred. A Zero Buffer Credit condition occurs when a port is unable to send frames because of a lack of buffer credit since the last node reset. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

Class 3 Send Timeout Frame Rate

The average number of class 3 frames per second that were discarded before transmission because of a timeout condition. The timeout condition occurs while the switch or port waits for buffer credit from the receiving port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

Class 3 Receive Timeout Frame Rate

The average number of class 3 frames per second that were discarded after reception because of a timeout condition. The timeout condition occurs while a transmitting port waits for buffer credit from a port at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

RDY Priority Override Rate

The average number of times per second during which the sending of R_RDY or VC_RDY signals was a higher priority than the sending of frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

Port State Change Rate

The average number of times per second that the state of a port changes to offline, online, or faulty. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

This property is available only for Brocade switches.

Port Congestion Index

The estimated degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally 0 - 100. The value 0 means there was no congestion. The value can exceed 100 if the buffer credit exhaustion persisted for an extended amount of time. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

Zero Buffer-to-Buffer Credit Percentage

The amount of time, as a percentage, that the port was not able to send frames between ports because of insufficient buffer-to-buffer credit. The amount of time value is measured from the last time that metadata was collected. In Fibre Channel technology, buffer-to-buffer credit is used to control the flow of frames between ports.

Credit Recovery Link Reset Rate

The estimated average number of link resets per second that a switch or port completed to recover buffer credits. This estimate attempts to disregard link resets that were caused by link initialization. When you troubleshoot a SAN, use this metric to help identify port conditions that might slow the performance of the resources to which those ports are connected.

Custom reports about capacity and relationships

You can use Cognos® Analytics to create custom reports about capacity and relationships. Capacity and relationship reports can contain detailed information about the properties of monitored resources, and the available space and capacity of those resources.

Capacity and relationship reports can contain information about the following resources and their components:

- Fabrics
- Switches
- Servers
- Hypervisors
- Network-attached storage servers
- Storage virtualizers
- Storage systems
- Storage resource groups
- Groups

In the IBM Spectrum® Control GUI, storage resource groups are available as general groups.

- [Creating custom capacity and relationship reports](#)
You can create custom reports that show capacity and relationship information for resources in the Cognos Analytics reporting tool. Capacity and relationship reports can contain detailed information about the properties of monitored resources, and the available space and capacity of those resources. You access custom capacity and relationship reports from IBM Spectrum Control, and you create the reports in the Cognos Analytics reporting tool.
- [Adding filters to reports](#)
You can use filters with capacity and relationship reports. Filters allow you to show only certain data that complies with your criteria. You add filters to your capacity and relationship reports when you create the reports.
- [Data and properties in capacity and relationship reports](#)
You can include capacity data, properties, relationship information, and other information about resources in your storage environment in capacity and relationship reports.

Creating custom capacity and relationship reports

You can create custom reports that show capacity and relationship information for resources in the Cognos® Analytics reporting tool. Capacity and relationship reports can contain detailed information about the properties of monitored resources, and the available space and capacity of those resources. You access custom capacity and relationship reports from IBM Spectrum® Control, and you create the reports in the Cognos Analytics reporting tool.

Before you begin


Before you can create a custom report about capacity and relationships, ensure that data was collected from the resources. The following table shows where the data must be collected from for each resource:

Resource	Prerequisite
Fabric	Data must be collected from the fabric by a CIM agent or SNMP agent.
Switch	Data must be collected from the switch by a CIM agent or SNMP agent.
Server	Data must be collected from the server by the Storage Resource agent, or by an agentless server.
Hypervisor	Data must be collected from the VMware ESX server.
NAS	Data must be collected from any NetApp or IBM® Scale Out Network Attached Storage resources that are in your storage environment.
Storage virtualizer	Data must be collected from the any of the following storage virtualizers that are in your storage environment: <ul style="list-style-type: none">• IBM FlashSystem® family• Storwize® V7000• Storwize V7000 Unified• SAN Volume Controller• Hitachi Universal Storage Platform V
Storage system	Data must be collected from the storage system, or any of the following systems that are configured as back-end storage: <ul style="list-style-type: none">• IBM FlashSystem family• Storwize V7000• Storwize V7000 Unified• SAN Volume Controller• Hitachi Universal Storage Platform V

Restriction: If you create a custom report that includes properties from both Servers to Storage Systems and Fabrics and Switches, the relationships between resources are not displayed correctly.

Tip: Before you run a report, you can preview the report with no data.

Procedure

1. Go to the URL for your Cognos Analytics server. The format of the URL is similar to this URL: `http://myhostname:9300/bi`
2. Depending on the type of report that you want to see, complete one of the following steps:
 - To create a basic report with limited formatting about the capacity and relationships of resources, complete the following steps:
 - Click New  in the Welcome portal.
 - Click Other, then click Query Studio.
 - Click IBM Spectrum Control Packages, then click the Capacity and Relationships package.
 - To create a report about the capacity and relationships of resources and apply advanced formatting features to the report, complete the following steps:
 - In the Welcome portal, click Team Content .
 - Click IBM Spectrum Control Packages.
 - Right-click Capacity and Relationships, then click Create report.
 - Click a template for the report.
3. Expand IBM Spectrum Control in the list of data sources.
4. Expand Servers to Storage Systems or Fabrics and Switches in the list.
5. Explore the folders in the list to locate the resources, properties, and statistics that you want to include in the report.
6. Depending on your browser, do one of the following steps to add properties and statistics to the report:
 - In Internet Explorer, drag the items that you want to see information about to the work area.
 - In Firefox, select the items that you want to see information about, and then click Insert.
7. Click Save.

Related tasks

- [Adding filters to reports](#)

Related information

-  [Run a report](#)

Adding filters to reports

You can use filters with capacity and relationship reports. Filters allow you to show only certain data that complies with your criteria. You add filters to your capacity and relationship reports when you create the reports.

About this task

Filters for resources are in the same list as the resources, in the appropriate resource folder. For example, to view the filters for servers, select Server Resources > Server Resource Filters > Server Filters.


Filters are available for all resources, except network-attached storage resources.

Procedure

1. To apply a filter to the data in the work area, locate the filter that you want to include in the folders in the list.
2. Depending on your browser, do one of the following steps to add a filter to the report:
 - In Internet Explorer, drag the filter that you want to use to the work area.
 - In Firefox, select the filter that you want to use, and then click Insert.The filter is applied to the report. If you did not add properties and statistics to the report, the filter is applied when you add the properties and statistics.
3. Optional: To delete the filter, right-click the filter in the work area, and then select Delete.

- **[Filters for servers in capacity and relationship reports](#)**
You can filter capacity and relationship reports to include only the servers, server file systems, and server logical volumes that meet specific criteria.
- **[Filters for hypervisors in capacity and relationship reports](#)**
You can filter capacity and relationship reports to include only the hypervisors, hypervisor file systems, and hypervisor logical volumes that meet specific criteria.
- **[Filters for storage virtualizers in capacity and relationship reports](#)**
You can filter capacity and relationship reports to include only the storage virtualizers and storage virtualizer volumes that meet specific criteria.
- **[Filters for storage systems in capacity and relationship reports](#)**
You can filter capacity and relationship reports to include only the storage systems or storage volumes that meet specific criteria.
- **[Filters for storage resource groups in capacity and relationship reports](#)**
You can filter capacity and relationship reports to include only the storage resource groups that meet specific criteria.

Related information

-  [Create a Simple Filter](#)

Filters for servers in capacity and relationship reports

You can filter capacity and relationship reports to include only the servers, server file systems, and server logical volumes that meet specific criteria.

Filters for servers

You can create capacity and relationship reports that include the following filters:

Servers Connected to Virtualizers or Storage Systems

Shows servers that are connected to storage virtualizers or storage systems. You can use this filter to identify servers that are connected directly to storage virtualizers, or connected directly or indirectly to storage systems. This filter excludes local disks from the report.

Servers Directly Connected to Storage Systems

Shows servers that are directly connected to a back-end storage system. You can use this filter to identify servers or server disks that are connected to a storage system, but are not connected through a storage virtualizer.

For example, you might want to migrate data from one storage system to another. In this case, you can use this filter to identify file systems or server disks that are not virtualized. These file systems or server disks might be affected if a storage system is removed.

Servers Directly Connected to Virtualizers

Shows servers that are directly connected to a storage virtualizer. You can use this filter to identify servers that are connected to a storage virtualizer, but are not connected directly to a back-end storage system.

For example, you can use this filter to identify only server disks that are on a SAN Volume Controller.

Filters for server file systems

You can create capacity and relationship reports that include the following filters:

Rows with File System Mount Points

Shows the file systems that have a mount point on the server.

Rows without File System Mount Points

Shows resources that do not have a mount point on the server. You can use this filter to identify disks that do not have file systems, that is, disks that are not in use. For example, to identify unused space include the Server Name property and the Server Disk Path property in a report, and then include the Rows without File System Mount Points filter.

Filters for server logical volumes

You can create capacity and relationship reports that include the following filters:

Rows with Logical Volume Path

Shows the logical volumes that have a path to the server.

Rows without Logical Volume Path

Shows resources that do not have a logical volume path to the server. You can use this filter to identify volumes that do not have paths, that is, volumes that are not in use.

For example, to identify unused space include the Server Name property and the Server Disk Path property in a report, and then include the Rows without Logical Volume Path filter.

Filters for hypervisors in capacity and relationship reports

You can filter capacity and relationship reports to include only the hypervisors, hypervisor file systems, and hypervisor logical volumes that meet specific criteria.

Filters for hypervisors

You can create capacity and relationship reports that include the following filters:

Hypervisor with Probed Servers

Shows hypervisors that are connected to servers. You can use this filter to include only hypervisors and their associated servers from which data was collected. For example, you might collect data only from your critical production servers on a hypervisor, but not your test servers. In this case, you can use this filter to exclude information for the test servers from the report.

Hypervisors Connected to Virtualizers or Storage Systems

Shows hypervisors that are connected to storage virtualizers or storage systems. You can use this filter to identify hypervisors that are connected directly to storage virtualizers, or connected directly or indirectly to storage systems. This filter excludes local disks from the report.

Hypervisors Directly Connected to Storage Systems

Shows hypervisors that are directly connected to a back-end storage system. You can use this filter to identify hypervisors or hypervisor disks that are connected to a storage system, but are not connected through a storage virtualizer.

For example, you might want to migrate data from one storage system to another. In this case, you can use this filter to identify file systems or hypervisor disks that are not virtualized. These file systems or hypervisor disks might be affected if a storage system is removed.

Hypervisors Directly Connected to Virtualizers

Shows hypervisors that are directly connected to a storage virtualizer. You can use this filter to identify hypervisors that are connected to a storage virtualizer, but are not connected directly to a back-end storage system.

For example, you can use this filter to identify only hypervisor disks that are on a SAN Volume Controller.

Filters for hypervisor file systems

You can create capacity and relationship reports that include the following filters:

Hypervisors with File System Mount Points

Shows the file systems that have a mount point on the hypervisor.

Hypervisors without File System Mount Points

Shows resources that do not have a mount point on the hypervisor. You can use this filter to identify disks that do not have file systems, that is, disks that are not in use.

For example, to identify unused space include the Hypervisor Name property and the Hypervisor Disk Path property in a report, and then include the Hypervisors without File System Mount Points filter.

Filters for hypervisor logical volumes

You can create capacity and relationship reports that include the following filters:

Hypervisors with Logical Volume Path

Shows the logical volumes that have a path to the hypervisor.

Hypervisors without Logical Volume Path

Shows resources that do not have a logical volume path to the hypervisor. You can use this filter to identify volumes that do not have paths, that is, volumes that are not in use.

For example, to identify unused space include the Hypervisor Name property and the Hypervisor Disk Path property in a report, and then include the Hypervisors without Logical Volume Path filter.

Filters for storage virtualizers in capacity and relationship reports

You can filter capacity and relationship reports to include only the storage virtualizers and storage virtualizer volumes that meet specific criteria.

Storage virtualizer filters

You can create capacity and relationship reports that include the following filters:

Virtualizers Not Connected to Servers

Shows storage virtualizers that are not connected to servers.

Filters for storage virtualizer volumes

You can create capacity and relationship reports that include the following filters:

Rows with Virtualizer Volumes

Shows storage virtualizer disks that are connected to the storage virtualizer that have an associated volume.

Rows without Virtualizer Volumes

Shows storage virtualizer disks that are connected to the storage virtualizer that do not have an associated volume.

Virtualizer Primary Volume Information

Shows primary volumes that are connected to the storage virtualizer. You can use this filter if you want to get a total of the used space on primary volumes in your storage environment.

Virtualizer Volume Copy Information

Shows copy volumes that are connected to the storage virtualizer. You can use this filter if you want to get a total of the used space on volume copies in your storage environment.

Filters for storage systems in capacity and relationship reports

You can filter capacity and relationship reports to include only the storage systems or storage volumes that meet specific criteria.

Filters for storage systems

You can create capacity and relationship reports that include the following filters:

Storage Systems Directly Connected to Servers

Shows back-end storage systems that are directly connected to a server. You can use this filter to identify storage systems that are connected to a server, but are not connected through a storage virtualizer.

For example, you might want to migrate data from one storage system to another. In this case, you can use this filter to identify file systems or server disks that are not virtualized. These file systems or server disks might be affected if a storage system is removed.

Storage Systems Connected to Servers through Virtualizers

Shows storage systems that are indirectly connected to a server through a storage virtualizer. You can use this filter to identify storage systems that are connected to a storage virtualizer, but are not connected directly to a server.

For example, you can use this filter to identify only server disks that are on a SAN Volume Controller.

Storage Systems not Connected to Servers

Shows storage systems that are not connected to a server either directly, or indirectly through a storage virtualizer.

Rows with Storage Volumes

Shows storage resources that have an associated storage volume.

Rows without Storage Volumes

Shows storage resources that do not have an associated storage volume.

Filters for storage resource groups in capacity and relationship reports

You can filter capacity and relationship reports to include only the storage resource groups that meet specific criteria.

Filters for storage resource groups

You can create capacity and relationship reports that include the following filters:

Storage Resource Group for Volumes

Shows only the volumes that are in a storage resource group.

Storage Resource Group for Pools

Shows only the pools that are in a storage resource group.

Storage Resource Group for Storage Systems

Shows only the storage systems that are in a storage resource group.

Storage Resource Group for Servers

Shows only the servers that are in a storage resource group.

Storage Resource Group for File Systems

Shows only the file systems that are in a storage resource group.

Data and properties in capacity and relationship reports

You can include capacity data, properties, relationship information, and other information about resources in your storage environment in capacity and relationship reports.

- [**Data for clusters in capacity and relationship reports**](#)
You can include general information about clusters in capacity and relationship reports.
- [**Data for servers in capacity and relationship reports**](#)
You can include general information, capacity data, properties of components, and other information about servers in capacity and relationship reports.
- [**Data for file systems on servers in capacity and relationship reports**](#)
You can include general information, capacity data, and status information about file systems on servers in capacity and relationship reports.
- [**Data for logical volumes on servers in capacity and relationship reports**](#)
You can include general information, capacity data, and status information about logical volumes on servers in capacity and relationship reports.
- [**Data for volume groups on servers in capacity and relationship reports**](#)
You can include general information and capacity data about volume groups on servers in capacity and relationship reports.
- [**Data for server disks in capacity and relationship reports**](#)
You can include general information, capacity data, properties of components, and other information about server disks in capacity and relationship reports.
- [**Data for server groups in capacity and relationship reports**](#)
You can include general information and properties of server groups in capacity and relationship reports.
- [**Data for server controllers in capacity and relationship reports**](#)
You can include general information about server controllers in capacity and relationship reports.
- [**Data for multipath drivers on servers in capacity and relationship reports**](#)
You can include name and version information about multipath drivers on servers in capacity and relationship reports.
- [**Data for file system groups on servers in capacity and relationship reports**](#)
You can include general information about file system groups on servers, and properties of file system groups in capacity and relationship reports.
- [**Data for hypervisors in capacity and relationship reports**](#)
You can include general information, capacity data, properties, and other information about hypervisors in capacity and relationship reports.

- [Data for clusters on hypervisors in capacity and relationship reports](#)
You can include general information about clusters on hypervisors in capacity and relationship reports.
- [Data for file systems on hypervisors in capacity and relationship reports](#)
You can include general information, capacity data, and other information about file systems on hypervisors in capacity and relationship reports.
- [Data for data stores on hypervisors in capacity and relationship reports](#)
You can include general information, capacity data, and other information about data stores on hypervisors in capacity and relationship reports.
- [Data for server disks on hypervisors in capacity and relationship reports](#)
You can include general information, capacity data, path information, and other information about server disks on hypervisors in capacity and relationship reports.
- [Data for hypervisor controllers in capacity and relationship reports](#)
You can include general information about hypervisor controllers in capacity and relationship reports.
- [Data for multipath drivers on hypervisors in capacity and relationship reports](#)
You can include name and version information about multipath drivers on hypervisors in capacity and relationship reports.
- [Data for virtual machines managed by hypervisors in capacity and relationship reports](#)
You can include general information, and file and disk information for virtual machines in capacity and relationship reports.
- [Data for network-attached storage systems in capacity and relationship reports](#)
You can include general information, capacity data, properties, and other information about network-attached storage (NAS) systems in capacity and relationship reports.
- [Data for file systems on network-attached storage systems in capacity and relationship reports](#)
You can include general information, capacity data, and other information about file systems on network-attached storage (NAS) systems in capacity and relationship reports.
- [Data for network-attached storage exports in capacity and relationship reports](#)
You can include general information and status information about network-attached storage (NAS) exports in capacity and relationship reports.
- [Data for logical volumes on network-attached storage systems in capacity and relationship reports](#)
You can include general information, capacity data, and other information about logical volumes on network-attached storage (NAS) systems in capacity and relationship reports.
- [Data for network-attached storage disks in capacity and relationship reports](#)
You can include general information, capacity data, properties, and other information about network-attached storage (NAS) disks in capacity and relationship reports.
- [Data for pools on Storwize V7000 Unified systems in capacity and relationship reports](#)
You can include general information and capacity data about pools on Storwize® V7000 Unified systems in capacity and relationship reports.
- [Data for Storwize V7000 Unified filesets in capacity and relationship reports](#)
You can include general information, capacity data, and other information about Storwize V7000 Unified filesets in capacity and relationship reports.
- [Data for storage virtualizers in capacity and relationship reports](#)
You can include general information, capacity data, properties, and other information about storage virtualizers in capacity and relationship reports.
- [Data for storage virtualizer pools in capacity and relationship reports](#)
You can include general information, capacity data, properties, and other information about storage virtualizer pools in capacity and relationship reports.
- [Data for storage virtualizer volumes in capacity and relationship reports](#)
You can include general information, capacity data, properties, information about copies of volumes, and other information about storage virtualizer volumes in capacity and relationship reports.
- [Data for managed disks on storage virtualizers in capacity and relationship reports](#)
You can include general information, capacity data, properties, and other information about managed disks on storage virtualizers in capacity and relationship reports.
- [Data for local disks on storage virtualizers in capacity and relationship reports](#)
You can include general information, capacity data, properties, and other information about local disks on storage virtualizers in capacity and relationship reports.
- [Data for storage virtualizer groups in capacity and relationship reports](#)
You can include the name of storage virtualizer groups and properties for storage virtualizer groups in capacity and relationship reports.
- [Data for storage systems in capacity and relationship reports](#)
You can include general information, capacity data, properties of components, and other information about storage systems in capacity and relationship reports.
- [Data for storage system pools in capacity and relationship reports](#)
You can include general information, capacity data, properties of components, and other information about storage system pools in capacity and relationship reports.
- [Data for storage system volumes in capacity and relationship reports](#)
You can include general information, capacity data, properties, information about copies of volumes, and other information about storage system volumes in capacity and relationship reports.
- [Data for storage system disks in capacity and relationship reports](#)
You can include properties, and other data about storage system disks in capacity and relationship reports.
- [Data for storage system groups in capacity and relationship reports](#)
You can include general information about, and properties of, storage system groups in capacity and relationship reports.
- [Data for managed disks on storage systems in capacity and relationship reports](#)
You can include general information, capacity data, properties, and other information about managed disks on storage systems in capacity and relationship reports.
- [Data for storage resource groups in capacity and relationship reports](#)
You can include general information about storage resource groups in capacity and relationship reports. In the IBM Spectrum® Control GUI, storage resource groups are available as general groups.
- [Data for groups in capacity and relationship reports](#)
You can include general information about groups in capacity and relationship reports. A group is a set of logically related volumes, file systems, and shares. For example, a group that represents a business critical application might include the volumes, file systems, and shares that provide storage to the application.
- [Data for switches in capacity and relationship reports](#)
You can include properties and other information about switches in capacity and relationship reports.
- [Data for switch ports in capacity and relationship reports](#)
You can include properties and other information about switch ports in capacity and relationship reports.
- [Data for fabrics in capacity and relationship reports](#)
You can include properties and other information about fabrics in capacity and relationship reports.

Data for clusters in capacity and relationship reports

You can include general information about clusters in capacity and relationship reports.

Information about clusters

You can create capacity and relationship reports that include the following information:

Cluster Name

The name that was assigned to the cluster when it was created.

Cluster Type

The type of the cluster. The cluster can be a Microsoft Cluster Server or a High-Availability Cluster Multiprocessing (HACMP) server.

Cluster IP Address

The IP address of the resource.

Cluster Domain Name

The domain name of the cluster.

Data for servers in capacity and relationship reports

You can include general information, capacity data, properties of components, and other information about servers in capacity and relationship reports.

Information about servers

You can create capacity and relationship reports that include the following information:

Server Name

The fully qualified domain name of the server. For example, the name of a server might be `server.example.com`.

Server Domain Name

The domain name of the server.

If a IBM Spectrum® Control server uses one DNS server, and another server, which has a Storage Resource agent installed, uses a different DNS server, the IBM Spectrum Control server might be displayed in some IBM Spectrum Control pages with different fully qualified host names. This is because the IBM Spectrum Control server is identified by two different DNS servers.

Server Short Name

The host name from the fully qualified domain name of the server. For example, if the fully qualified domain name of the server is `mycomputer.example.com`, the host name is `mycomputer`.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Server Swap Space (GiB)

The amount of disk space on a resource that is available to store some of the contents of the RAM when the RAM is full.

Server Total Disk Capacity (GiB)

The amount of storage space that is on a server.

Server Total Disk Available Space (GiB)

The amount of unused storage space on all disks on a server.

Server Total File System Available Space (GiB)

The amount of unused storage space that is available on the file systems on a server.

Component properties

You can create capacity and relationship reports that include the following information:

Server Cluster Current Hosting Node

The name of the server that is hosting the cluster resource group.

Server Cluster Resource Type

Shows whether the cluster is a node or a cluster resource group.

Server IP Address

The IP address of the resource.

Server Is Virtual Machine

Shows whether a resource is a virtual machine. If this value is **Yes**, the resource is a virtual machine.

Server OS Type

The operating system that is running on the server.

Server OS Version

The version number of the operating system that is running on the resource.

Server Processor Architecture

The architecture of the processor on the server or hypervisor. For example, the architecture of a processor might be Intel 64 bit (IA64) or Intel 32 bit (IA32).

Server Processor Count

The number of processors on the server or hypervisor.

Server Processor Speed (MHz)

The speed of the processor on the server or hypervisor.

Server Processor Type

Shows information about the processor, such as the family and model of the processor.

Server RAM (GiB)

The amount of RAM on the server or hypervisor.

Server Time Zone

The time zone in which a resource is located.

Server Location

The physical location of a server. The location is defined when a server is added to IBM Spectrum Control. You can add or edit the location of the server in the Properties pane of the server.

Server Custom Tag 1, 2, and 3

User-defined text that is associated with a server. You can add or edit the custom tags for a server in the Properties pane of the server.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Server Model

The model name or model number of the resource.

Server Serial Number

The serial number of the resource.

Server Vendor

The vendor who supplied the resource.

Status information

You can create capacity and relationship reports that include the following information:

Server Status

The status of a server, such as **Normal**, **Warning**, **Error**, **Unreachable**, **Unknown**, or **Agentless**. Use the status to determine the condition of the server, and if any actions must be taken. If this value is **Agentless**, the server is an unmanaged server.

Server Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Server Last Data Collection

The date and time when storage statistics were last collected from the resource.

Server Last Start Time

The last time that the resource was started.

Data for file systems on servers in capacity and relationship reports

You can include general information, capacity data, and status information about file systems on servers in capacity and relationship reports.

Information about file systems on servers

You can create capacity and relationship reports that include the following information:

File System Mount Point

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be c:\ or d:\. On operating systems such as AIX®, Linux®, Solaris, or HP-UX, the mount point might be /opt or /export/home.

File System Export Name

The name of the exported file system.

File System Type

The type of file system that the resource uses.

Capacity and usage data

Tip:

In the IBM Spectrum® Control interface, the capacity values of file systems are displayed in one column, Total File System Capacity (GiB).

You can create capacity and relationship reports that include the following information:

File System Capacity (GiB)

The amount of storage space on the file system of the resource.

File System Directory Count

The number of directories in the file systems that are on a server.

File System File Count

The number of files in the file system.

File System Available Space (GiB)

The amount of unused storage space in the file system of the resource.

File System Available Inodes

The number of unused inodes in file systems on the operating system.

File System Used Inodes

The number of used inodes in file systems on the operating system.

Tip: For Microsoft Windows systems, this property is blank.

File System Maximum Files

The maximum number of files that the file system on the resource can contain.

File System Physical Size (GiB)

The amount of physical storage space on the file systems on a resource. The physical size is the size of all the clusters that the file system uses.

File System Used Space (GiB)

The amount of used storage space in the file system of the resource.

Component properties

You can create capacity and relationship reports that include the following information:

File System Is Remote Mount

Shows whether the file system is remotely mounted on another server. If this value is **Yes**, the file system is remotely mounted.

Status information

You can create capacity and relationship reports that include the following information:

File System Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

File System Last Scan

The date and time when file statistics were last collected from the resource.

Data for logical volumes on servers in capacity and relationship reports

You can include general information, capacity data, and status information about logical volumes on servers in capacity and relationship reports.

Information about logical volumes on servers

You can create capacity and relationship reports that include the following information:

Logical Volume Path

The path to a logical volume on a resource, for example `/dev/hd1`.

Logical Volume Is Swap Space

Shows whether a logical volume is used to store some of the contents of the RAM when the RAM is full. If this value is **Yes**, the logical volume is used to store some of the contents of the RAM.

Logical Volume Mirror Count

The number of mirrors that keep a synchronous copy of the resource.

Logical Volume Type

The type of server logical volume. For example, the logical volume type can be concatenation or stripe set.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Logical Volume Allocation on Server Disk (GiB)

The amount of disk storage space that is allocated to the logical volume.

Logical Volume Capacity (GiB)

The amount of storage space that is on a logical volume.

Logical Volume Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Status information

You can create capacity and relationship reports that include the following information:

Logical Volume Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Data for volume groups on servers in capacity and relationship reports

You can include general information and capacity data about volume groups on servers in capacity and relationship reports.

Information about volume groups on servers

You can create capacity and relationship reports that include the following information:

Volume Group Name

The name that was assigned to the volume group when it was added to the system.

Volume Group Type

The type of the volume manager that manages the volume group. The volume manager can be a logical volume manager (LVM) or a Veritas volume manager.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Volume Group Capacity (GiB)

The amount of storage space that is on a volume group.

Volume Group Available Space (GiB)

The amount of unused storage space on a server volume group.

Volume Group Number of Logical Volumes

The number of logical volumes on a resource.
Volume Group Number of Server Disks
The number of server disks in a volume group.

Data for server disks in capacity and relationship reports

You can include general information, capacity data, properties of components, and other information about server disks in capacity and relationship reports.

Information about server disks

You can create capacity and relationship reports that include the following information:

Server Disk Name
The path that the operating system uses for the server disk, for example /dev/hdisk0.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Server Disk Capacity (GiB)
The amount of storage space that is on a server disk.
Server Disk Available Space (GiB)
The amount of unused storage space that is on a server disk.

Component properties

You can create capacity and relationship reports that include the following information:

Server Disk Is Removable
Shows whether the disk can be removed from the server. If this value is **Yes**, the disk can be removed from the server.
Server Disk Multipathing Policy
The multipathing policy that is in effect for a disk. For example, the policy can be **Round Robin**, **Load Balancing**, **Failover Only**, or other policies.

Status information

You can create capacity and relationship reports that include the following information:

Server Disk Status
The condition of the resource, for example normal, warning, or error.
Server Disk Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Server Disk Firmware Version
The version number of the firmware that is running on the disk.
Server Disk Model
The model name or model number of the resource.
Server Disk Serial Number
The serial number of the resource.
Server Disk Vendor
The vendor who supplied the resource.

Path information

You can create capacity and relationship reports that include the following information:

Server Disk Path Name
The name of a path to a server disk. Typically the multipath driver defines this name.
Server Disk Path Mode
The online status of the path to the server disk.
Server Disk Path State
The state of the disk path shows whether the path is functioning correctly.
Server Disk Path Status
The aggregate status of the paths to the server disk.
Server Disk Preferred Path
The preferred path status of the path to the server disk, that is, whether the path is the preferred path for the disk. The multipath driver might use non-preferred paths if a failover occurs. The disk might have no preferred path, depending on the multipathing policy that is in effect for the disk.
Server Disk LUN
The logical unit within a small computer system interface target on a disk.
Server Disk SCSI Target
The bus address for the small computer system interface of the target on a disk.

Server Disk Target WWPN
The worldwide port name of the storage port that the volume is mapped to.

Data for server groups in capacity and relationship reports

You can include general information and properties of server groups in capacity and relationship reports.

Information about server groups

Server groups are deprecated in IBM Spectrum® Control 5.2.10 and later. Use storage resource groups instead. In the IBM Spectrum Control GUI, storage resource groups are available as general groups.

You can create capacity and relationship reports that include the following information:

Server Group Name
The name that was assigned to the server group when it was added to the system.

Component properties

You can create capacity and relationship reports that include the following information:

Server Group Creator
The user name of the user who created the group.
Server Group Last Modified
The date and time that the resource was last modified.
Server Group Last Modified by User
The user name of the user who last modified the resource.

Data for server controllers in capacity and relationship reports

You can include general information about server controllers in capacity and relationship reports.

Information about server controllers

You can create capacity and relationship reports that include the following information:

Server Controller Name
The name of the disk controller on a server.
Server Controller Status
The status of a disk controller. Statuses include Normal, Warning, Error, and Unknown. Use the status to determine the condition of a controller, and if any actions must be taken. For example, if a controller has an Error status, take immediate action to correct the problem.
Server Controller Description
A short description of a disk controller that was provided by the manufacturer.
Server Controller Type
The type of disk controller, such as IDE, SCSI, Floppy, and RAID. Host Bus Adapters (HBAs) have a controller type of Fibre Channel Arbitrated Loop (FCAL).
Server Controller Driver Version
The version identifier of the device driver on the disk controller. This value is only available for HBAs.
Server Controller Firmware
The firmware version of the microcode on a disk controller. This value is only available for HBAs.
Server Controller ROM Version
The version of the read-only memory (ROM) on a controller. This value is only available for HBAs.
Server Controller Hardware Version
The hardware version identifier for a disk controller. This value is only available for HBAs.
Server Controller WWN
The worldwide name of a disk controller. This value is only available for HBAs.
Server Controller Serial Number
The serial number of the server controller.
Server Controller Bus Address
The SCSI bus address that is associated with a disk controller. This property is available only for SCSI disk controllers.
Server Controller Bus Number
The SCSI bus number that is associated with a disk controller. This property is available only for SCSI disk controllers.

Data for multipath drivers on servers in capacity and relationship reports

You can include name and version information about multipath drivers on servers in capacity and relationship reports.

Information about multipath drivers on servers

You can create capacity and relationship reports that include the following information:

Server Multipath Driver Name
The name of a multipath driver on a server.

Server Multipath Driver Version
The version number of a multipath driver on a server.

Data for file system groups on servers in capacity and relationship reports

You can include general information about file system groups on servers, and properties of file system groups in capacity and relationship reports.

Information about file system groups on servers

File system groups are deprecated in IBM Spectrum® Control 5.2.10 and later. Use storage resource groups instead. In the IBM Spectrum Control GUI, storage resource groups are available as general groups.

You can create capacity and relationship reports that include the following information:

File System Group Name
The name that was assigned to the file system group when it was added to the system.

Component properties

You can create capacity and relationship reports that include the following information:

File System Group Creator
The user name of the user who created the group.

File System Group Last Modified
The date and time that the resource was last modified.

File System Group Last Modified by User
The user name of the user who last modified the resource.

Data for hypervisors in capacity and relationship reports

You can include general information, capacity data, properties, and other information about hypervisors in capacity and relationship reports.

Information about hypervisors

You can create capacity and relationship reports that include the following information:

Hypervisor Name
The fully qualified domain name of the hypervisor. For example, the name of a hypervisor might be **hypervisor.example.com**.

Hypervisor Short Name
The host name from the fully qualified domain name of the hypervisor. For example, if the fully qualified domain name of the hypervisor is **mycomputer.example.com**, the host name is **mycomputer**.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Hypervisor Kernel Memory (GiB)
The amount of memory on a hypervisor that is used for kernel memory. Kernel memory is used to perform internal operating system processes.

Hypervisor Service Console Memory (GiB)
The amount of memory on a hypervisor that is used for the VMware ESX Service Console. Service Console is the operating system that is used to interact with ESX and the virtual machines that run on the hypervisor.

Hypervisor Swap Space (GiB)
The amount of disk space on a resource that is available to store some of the contents of the RAM when the RAM is full.

Hypervisor Total Disk Capacity (GiB)
The amount of storage space on disks on a hypervisor.

Hypervisor Total Disk Available Space (GiB)
The amount of unused storage space on disks on the hypervisor.

Hypervisor Total File System Available Space (GiB)
The amount of unused storage space that is available on the file systems on a hypervisor.

Component properties

You can create capacity and relationship reports that include the following information:

Hypervisor IP Address
The IP address of the resource.

Hypervisor OS Type
The software that runs the hypervisor, that is, VMware ESX.

Hypervisor OS Version
The version number of the operating system that is running on the resource.

Hypervisor Processor Architecture

The architecture of the processor on the server or hypervisor. For example, the architecture of a processor might be Intel 64 bit (IA64) or Intel 32 bit (IA32).

Hypervisor Processor Count
The number of processors on the server or hypervisor.

Hypervisor Processor Speed (MHz)
The speed of the processor on the server or hypervisor.

Hypervisor Processor Type
Shows information about the processor, such as the family and model of the processor.

Hypervisor RAM (GiB)
The amount of RAM on the server or hypervisor.

Hypervisor Software API Version
The version number of the software API that runs on the hypervisor.

Hypervisor Software Full Name
The name, version number, and build number of the software that runs the hypervisor.

Hypervisor Time Zone
The time zone in which a resource is located.

Hypervisor Location
The physical location of a hypervisor. The location is defined when a hypervisor is added to IBM Spectrum® Control. You can add or edit the location of the hypervisor in the Properties pane of the hypervisor.

Hypervisor Custom Tag 1, 2, and 3
User-defined text that is associated with a hypervisor. You can add or edit the custom tags for a hypervisor in the Properties pane of the hypervisor.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Hypervisor Model
The model name or model number of the resource.

Hypervisor Serial Number
The serial number of the resource.

Hypervisor Vendor
The vendor who supplied the resource.

Status information

You can create capacity and relationship reports that include the following information:

Hypervisor Status
The condition of the resource, for example normal, warning, or error.

Hypervisor Is Detected
Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Hypervisor Last Data Collection
The date and time when storage statistics were last collected from the resource.

Hypervisor Last Start Time
The last time that the resource was started.

Data for clusters on hypervisors in capacity and relationship reports

You can include general information about clusters on hypervisors in capacity and relationship reports.

Information about clusters on hypervisors

You can create capacity and relationship reports that include the following information:

Hypervisor Cluster Name
The name of a cluster that is monitored in your storage environment. A cluster is a group of hypervisors that collaborate for the purposes of workload balancing and failover.

Hypervisor Cluster Type
The type of a cluster that is monitored in your storage environment. For example, the type might be VMware Cluster.

Data for file systems on hypervisors in capacity and relationship reports

You can include general information, capacity data, and other information about file systems on hypervisors in capacity and relationship reports.

Information about file systems on hypervisors

You can create capacity and relationship reports that include the following information:

Hypervisor File System Mount Point
The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be c:\ or d:\. On operating systems such as AIX®, Linux®, Solaris, or HP-UX, the mount point might be /opt or /export/home.

Hypervisor File System Export Name
The name of the exported file system.

Hypervisor File System Type

The type of file system that the hypervisor uses.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Hypervisor File System Capacity (GiB)

The amount of storage space on the file system of the resource.

Hypervisor File System Directory Count

The number of directories in the file systems that are on a server.

Hypervisor File System File Count

The number of files in the file system.

Hypervisor File System Available Space (GiB)

The amount of unused storage space in the file system of the resource.

Hypervisor File System Available Inodes

The number of unused inodes in file systems on the operating system.

Hypervisor File System Used Inodes

The number of used inodes in file systems on the operating system.

Tip: For Microsoft Windows systems, this property is blank.

Hypervisor File System Maximum Files

The maximum number of files that the file system on the resource can contain.

Hypervisor File System Physical Size (GiB)

The amount of physical storage space on the file systems on a resource. The physical size is the size of all the clusters that the file system uses.

Hypervisor File System Used Space (GiB)

The amount of used storage space in the file system of the resource.

Status information

You can create capacity and relationship reports that include the following information:

Hypervisor File System Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Hypervisor File System Last Scan

The date and time when file statistics were last collected from the resource.

Data for data stores on hypervisors in capacity and relationship reports

You can include general information, capacity data, and other information about data stores on hypervisors in capacity and relationship reports.

Information about data stores on hypervisors

You can create capacity and relationship reports that include the following information:

Hypervisor Data Store

The path to the VMware ESX data store on the hypervisor, for example datastore_svc3c.

Hypervisor Data Store Cluster Name

The name of the cluster, if the data store is a member of a data store cluster. If the data store is not a member of a data store cluster, no information is displayed.

Hypervisor Data Store Type

The type of file system that the hypervisor data store uses. For example, the type can be NFS or VMFS.

Hypervisor NAS Name

The name of the network-attached storage that the data store is obtained from.

Hypervisor NAS Path Name

The path name on the network-attached storage of the data store.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Hypervisor Data Store Capacity (GiB)

The amount of storage space that is on the hypervisor data store.

Hypervisor Data Store Available Space (GiB)

The amount of unused storage space on a hypervisor data store.

Hypervisor Data Store Allocation on Hypervisor Disk (GiB)

The amount of storage space that is allocated to the hypervisor data store on hypervisor server disks.

Hypervisor Data Store VM File System Version

The version number of the virtual machine file system on the hypervisor data store.

Status information

You can create capacity and relationship reports that include the following information:

Hypervisor Data Store Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Data for server disks on hypervisors in capacity and relationship reports

You can include general information, capacity data, path information, and other information about server disks on hypervisors in capacity and relationship reports.

Information about server disks

You can create capacity and relationship reports that include the following information:

Hypervisor Disk Name

The path that the operating system uses for the server disk.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Hypervisor Disk Capacity (GiB)

The amount of storage space that is on the hypervisor disk.

Hypervisor Disk Available Space (GiB)

The amount of unused storage space on a hypervisor server disk.

Component properties

You can create capacity and relationship reports that include the following information:

Hypervisor Disk Multipathing Policy

The multipathing policy that is in effect for a disk. For example, the policy can be **Round Robin, Load Balancing, Failover Only**, or other policies.

Status information

You can create capacity and relationship reports that include the following information:

Hypervisor Disk Status

The condition of the resource, for example normal, warning, or error.

Hypervisor Disk Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Hypervisor Disk Firmware Version

The version number of the firmware that is running on the disk.

Hypervisor Disk Model

The model name or model number of the resource.

Hypervisor Disk Serial Number

The serial number of the resource.

Hypervisor Disk Vendor

The vendor who supplied the resource.

Path information

You can create capacity and relationship reports that include the following information:

Hypervisor Disk Path Name

The name of a path to a server disk. Typically the multipath driver defines this name.

Hypervisor Disk Path State

The state of the disk path shows whether the path is functioning correctly.

Hypervisor Disk Path Status

The aggregate status of the paths to the server disk.

Hypervisor Disk Preferred Path

The preferred path status of the path to the server disk, that is, whether the path is the preferred path for the disk. The multipath driver might use non-preferred paths if a failover occurs. The disk might have no preferred path, depending on the multipathing policy that is in effect for the disk.

Hypervisor Disk LUN

The logical unit within a small computer system interface target on a disk.

Hypervisor Disk SCSI Target

The bus address for the small computer system interface of the target on a disk.

Hypervisor Disk Target WWPN

The worldwide port name of the storage port that the volume is mapped to.

Data for hypervisor controllers in capacity and relationship reports

You can include general information about hypervisor controllers in capacity and relationship reports.

Information about hypervisor controllers

You can create capacity and relationship reports that include the following information:

Hypervisor Controller Name

The name of the disk controller on a hypervisor.

Hypervisor Controller Description

A short description of a disk controller that was provided by the manufacturer.

Hypervisor Controller Type

The type of disk controller, such as IDE, SCSI, Floppy, and RAID. Host Bus Adapters (HBAs) have a controller type of Fibre Channel Arbitrated Loop (FCAL).

Hypervisor Controller WWN

The worldwide name of a disk controller. This value is only available for HBAs.

Hypervisor Controller Bus Address

The SCSI bus address that is associated with a disk controller. This property is available only for SCSI disk controllers.

Data for multipath drivers on hypervisors in capacity and relationship reports

You can include name and version information about multipath drivers on hypervisors in capacity and relationship reports.

Information about multipath drivers on hypervisors

You can create capacity and relationship reports that include the following information:

Hypervisor Multipath Driver Name

The name of a multipath driver on a hypervisor.

Hypervisor Multipath Driver Version

The version number of a multipath driver on a hypervisor.

Data for virtual machines managed by hypervisors in capacity and relationship reports

You can include general information, and file and disk information for virtual machines in capacity and relationship reports.

Information about virtual machines

You can create capacity and relationship reports that include the following information:

Virtual Machine Name

The name that was assigned to the virtual machine when it was added to the system.

Virtual Machine Configuration File

The path to the virtual machine configuration file for a hypervisor virtual machine, for example [datastore1] example/example.vmx.

Virtual Machine Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Virtual Machine Number of CPU

The number of processors on the virtual machine.

Virtual Machine Status

The condition of the resource, for example normal, warning, or error.

Virtual Machine Total Memory (GiB)

The amount of memory on a hypervisor virtual machine.

Virtual machine file information

You can create capacity and relationship reports that include the following information:

Virtual Machine File Path

The path to the virtual machine configuration file, or to the virtual machine disk file for the virtual machine. The virtual machine configuration file is a .vmx file, and the virtual machine disk file is a .vmdk file.

Virtual Machine File Logical Size (GiB)

The logical size of a virtual machine disk file for a hypervisor virtual machine.

Virtual Machine File Type

The type of the file for the virtual machine. For example, the file can be a virtual machine configuration file, a virtual machine disk file, an ISO image file, or other types of virtual machine file.

Virtual machine disk information

You can create capacity and relationship reports that include the following information:

Virtual Machine Disk Name

The name that was assigned to the virtual machine disk when it was added to the system.

Virtual Machine Disk SCSI Target

The bus address for the small computer system interface of the target on a disk.

Virtual Machine Disk Type

The type of virtual machine disk, for example, collection mirroring or logical disk mirroring.

Data for network-attached storage systems in capacity and relationship reports

You can include general information, capacity data, properties, and other information about network-attached storage (NAS) systems in capacity and relationship reports.

Information about NAS systems

You can create capacity and relationship reports that include the following information:

NAS Name

The fully qualified domain name of the NAS device. For example, the name of a NAS device might be **nas.example.com**.

NAS Domain Name

The domain name of the cluster, server, or NAS server.

NAS Short Name

The host name from the fully qualified domain name of the NAS device. For example, if the fully qualified domain name of the NAS device is **mycomputer.example.com**, the host name is **mycomputer**.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

NAS Total Disk Capacity (GiB)

The amount of storage space on all disks on a NAS resource.

NAS Total Disk Available Space (GiB)

The amount of unused storage space on all disks on a NAS resource.

NAS Total File System Available Space (GiB)

The amount of unused storage space on the file systems on a NAS resource.

Component properties

You can create capacity and relationship reports that include the following information:

NAS OS Type

The operating system that is running on the server.

NAS OS Version

The version number of the operating system that is running on the resource.

NAS Location

The physical location of a NAS system that is probed as a file server. You can add or edit the location of the NAS system on the File Storage tab in the Storage Systems pane.

NAS Custom Tag 1, 2, and 3

User-defined text that is associated with a NAS system that is probed as a file server. You can add or edit the custom tags for a NAS system on the File Storage tab in the Storage Systems pane.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

NAS Model

The model name or model number of the resource.

NAS Vendor

The vendor who supplied the resource.

NAS Serial Number

The serial number of the resource.

Status information

You can create capacity and relationship reports that include the following information:

NAS Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

NAS Last Data Collection

The date and time when storage statistics were last collected from the resource.

NAS Last Start Time

The last time that the resource was started.

NAS system components information

You can create capacity and relationship reports that include the following information:

NAS System Name

The name that was assigned to the NAS system when it was added to the system.

NAS System IP Address

The IP address of the resource.

NAS System Location

The physical location of a NAS system that is probed as a storage system. The location is defined when a NAS system is added to IBM Spectrum® Control. You can add or edit the location of the NAS system on the File Storage tab in the Storage Systems pane.

NAS System Custom Tag 1, 2, and 3

User-defined text that is associated with a NAS system that is probed as a storage system. You can add or edit the custom tags for a NAS system on the File Storage tab in the Storage Systems pane.

Data for file systems on network-attached storage systems in capacity and relationship reports

You can include general information, capacity data, and other information about file systems on network-attached storage (NAS) systems in capacity and relationship reports.

Information about file systems on NAS systems

You can create capacity and relationship reports that include the following information:

NAS File System Mount Point

The name or mount point of the file system for the resource. For example, on Microsoft Windows systems, the value in this property might be c:\ or d:\. On operating systems such as AIX®, Linux®, Solaris, or HP-UX, the mount point might be /opt or /export/home.

NAS File System Export Name

The name of the exported file system.

NAS File System Type

The type of file system that the resource uses.

Capacity and usage data

Tip:

In the IBM Spectrum® Control interface, the capacity values of file systems are displayed in one column, Total File System Capacity (GiB).

You can create capacity and relationship reports that include the following information:

NAS File System Capacity (GiB)

The amount of storage space on the file system of the resource.

NAS File System Directory Count

The number of directories in the file systems that are on a server.

NAS File System File Count

The number of files in the file system.

NAS File System Available Space (GiB)

The amount of unused storage space in the file system of the resource.

NAS File System Available Inodes

The number of unused inodes on the file system on a NAS server.

NAS File System Used Inodes

The number of used inodes on the file system on a NAS server.

NAS File System Maximum Files

The maximum number of files that the file system on the resource can contain.

NAS File System Physical Size (GiB)

The amount of physical storage space on the file systems on a resource. The physical size is the size of all the clusters that the file system uses.

NAS File System Used Space (GiB)

The amount of used storage space in the file system of the resource.

Component properties

You can create capacity and relationship reports that include the following information:

NAS File System Custom Tag 1, 2, and 3

User-defined text that is associated with a NAS file system. You can add or edit the custom tags for a NAS file system in the Properties pane of the resource.

Status information

You can create capacity and relationship reports that include the following information:

NAS File System Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

NAS File System Last Scan

The date and time when file statistics were last collected from the resource.

Data for network-attached storage exports in capacity and relationship reports

You can include general information and status information about network-attached storage (NAS) exports in capacity and relationship reports.

Information about network-attached storage exports

You can create capacity and relationship reports that include the following information:

NAS Export Name

The name of the network-attached storage export. Use this name to mount the network-attached storage export on a server.

NAS Export Path

The path to a logical volume on a resource, for example /dev/hd1.

NAS Export State

Shows whether the network-attached storage export is active.

Status information

You can create capacity and relationship reports that include the following information:

NAS Export Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Data for logical volumes on network-attached storage systems in capacity and relationship reports

You can include general information, capacity data, and other information about logical volumes on network-attached storage (NAS) systems in capacity and relationship reports.

Information about logical volumes on NAS systems

You can create capacity and relationship reports that include the following information:

NAS Logical Volume Name

The name that was assigned to the NAS logical volume when it was added to the system.

NAS Logical Volume Type

The type of NAS logical volume. The logical volume type can be a NetApp traditional volume or a NetApp flexible volume.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

NAS Logical Volume Capacity (GiB)

The amount of storage space on the NAS logical volume.

NAS Logical Volume Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Status information

You can create capacity and relationship reports that include the following information:

NAS Logical Volume Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

NAS Logical Volume Last Scan

The date and time when file statistics were last collected from the resource.

Data for network-attached storage disks in capacity and relationship reports

You can include general information, capacity data, properties, and other information about network-attached storage (NAS) disks in capacity and relationship reports.

Information about NAS disks

You can create capacity and relationship reports that include the following information:

NAS Disk Name

The name that was assigned to the NAS disk when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

NAS Disk Total Capacity (GiB)

The amount of storage space that is on the NAS disk.

NAS Disk Available Space (GiB)

The amount of unused storage space on a NAS disk.

Status information

You can create capacity and relationship reports that include the following information:

NAS Disk Status

The condition of the resource, for example normal, warning, or error.

NAS Disk Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Storwize® V7000 Unified information

You can create capacity and relationship reports that include the following information:

NSD Available Fragments

The number of fragments of disk space that are unused on the Network Shared Disk (NSD).

NSD Failover Group

The failover group that the NSD belongs to. A failover group is a group of resources that are grouped together to ensure that some of those resources are available if one or more of the resources fail.

NSD Type

The type of the NSD. For example, the type can be:

- Data and metadata
- Metadata only
- Data only

NSD Last Data Collection

The date and time when storage statistics were last collected from the resource.

NSD Custom Tag 1, 2, and 3

User-defined text that is associated with a Network Shared Disk (NSD). You can add or edit the custom tags for an NSD in the Properties pane of the resource.

NetApp information

You can create capacity and relationship reports that include the following information:

NAS Disk Flag

Shows whether the NAS disk is a spare device. If this value is **Spare device**, the NAS disk is a spare device.

NAS Disk Is Encrypted

Shows whether the resource is encrypted. If this value is **Yes**, the resource is encrypted.

NAS Disk Is Removable

Shows whether the disk can be removed from the server. If this value is **Yes**, the disk can be removed from the server.

NAS Disk Is Solid State

Shows whether the disk is a solid-state drive.

NetApp paths information

You can create capacity and relationship reports that include the following information:

NAS Disk LUN

The logical unit within a small computer system interface target on a disk.

NAS Disk SCSI Target

The bus address for the small computer system interface of the target on a disk.

NAS Disk Target WWPN

The worldwide port name of the storage port that the volume is mapped to.

NetApp vendor, model, and device information

You can create capacity and relationship reports that include the following information:

NAS Disk Firmware Version

The version number of the firmware that is running on the disk.

NAS Disk Model

The model name or model number of the resource.

NAS Disk Serial Number

The serial number of the resource.

NAS Disk Vendor

The vendor who supplied the resource.

Data for pools on Storwize V7000 Unified systems in capacity and relationship reports

You can include general information and capacity data about pools on Storwize® V7000 Unified systems in capacity and relationship reports.

Information about pools on network-attached storage systems

You can create capacity and relationship reports that include the following information:

NAS Pool Name

The name that was assigned to the NAS pool when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

NAS Pool Capacity (GiB)

The amount of storage space that is in a NAS pool.

NAS Pool Available Space (GiB)

The amount of unused storage space in a NAS pool.

NAS Pool Number of Logical Volumes

The number of logical volumes on a resource.

Data for Storwize V7000 Unified filesets in capacity and relationship reports

You can include general information, capacity data, and other information about Storwize® V7000 Unified filesets in capacity and relationship reports.

Information about network-attached storage filesets

You can create capacity and relationship reports that include the following information:

NAS Fileset Name

The name that was assigned to the network-attached storage (NAS) fileset when it was added to the system.

NAS Fileset Comment

A comment that is associated with the NAS fileset.

NAS Fileset Path

The path to a logical volume on a resource, for example /dev/hd1.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

NAS Fileset Used Space (GiB)

The amount of used storage space in the fileset of the NAS resource.

NAS Fileset Used Inodes

The number of used inodes in file systems on operating systems such as AIX®, Linux®, Solaris, or HP-UX.

Status information

You can create capacity and relationship reports that include the following information:

NAS Fileset Status

Shows whether the network-attached storage fileset is linked or unlinked.

NAS Fileset Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Data for storage virtualizers in capacity and relationship reports

You can include general information, capacity data, properties, and other information about storage virtualizers in capacity and relationship reports.

Information about storage virtualizers

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Name

A user-defined name of the storage virtualizer. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage virtualizer was added for monitoring.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Total Pool Capacity (GiB)

The total amount of storage space in pools on the storage virtualizer.

Storage Virtualizer Pool Used Space (GiB)

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For resources that are running IBM Spectrum Virtualize, note the following information about the Storage Virtualizer Pool Used Space:

- You can pre-allocate thin-provisioned volume space when the volumes are created. In this case, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes.

- For pools with compressed volumes, the Storage Virtualizer Pool Used Space reflects the size of compressed data that is written to disk. As the data changes, the Storage Virtualizer Pool Used Space might at times be less than the Storage Virtualizer Pool Allocated Space.
- For pools with volumes that are not thin provisioned or compressed, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that IBM Spectrum Control collected data about a pool. Because data collection is run on a set schedule and the used space on volumes can change rapidly, the value in this column might not be 100% accurate for the current state of volumes.

Storage Virtualizer Pool Available Space (GiB)

The amount of space in a pool that is not reserved for volumes.

Storage Virtualizer Total Volume Capacity (GiB)

The amount of virtual volume space on all volumes on the storage virtualizer.

Storage Virtualizer Volume Allocated Space (GiB)

The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following resources:

- Resources that are not running IBM Spectrum Virtualize
- Resources that are running IBM Spectrum Virtualize and are not thin-provisioned

Storage Virtualizer Volume Assigned Space (GiB)

The amount of virtual volume space that is assigned to the server or storage virtualizer.

Storage Virtualizer MDisk Available Space (GiB)

The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

Storage Virtualizer MDisk Capacity (GiB)

The amount of storage space on the managed disk.

Storage Virtualizer Disk Capacity (GiB)

The amount of local disk storage on the storage virtualizer. The local disks are the internal disk drives.

Storage Virtualizer Number of MDisk

The number of managed disks on a resource.

Storage Virtualizer Number of Volumes

The number of volumes on a resource.

Storage Virtualizer Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Virtualizer Tier Available Space (GiB)

The amount of unused tier storage of a particular type. The type is identified in the Storage Virtualizer Tier Type property.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Tier Capacity (GiB)

The amount of tier storage of a particular type. The type is identified in the Storage Virtualizer Tier Type property.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Physical Allocation Percentage

The percentage of physical space in storage virtualizer pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

IBM Spectrum Control uses the following formula to determine the allocation percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total storage virtualizer pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Virtual Allocation Percentage

The percentage of physical space in storage virtualizer pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB.

This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage Virtualizer Shortfall Percentage property to estimate how critical the shortage of space is for storage virtualizer pools.

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Shortfall Percentage

The percentage of the remaining unallocated volume space in storage virtualizer pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{volume space} - \text{used volume space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space.

This Storage Virtualizer Shortfall Percentage represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

This property is available only for storage virtualizer pools with thin-provisioned volumes.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Access Information

The URL from which you can access the storage system or storage virtualizer.

Storage Virtualizer Alias

Shows an alternative name for a storage virtualizer.

Storage Virtualizer Code Level

The microcode or firmware level of the storage virtualizer.

Storage Virtualizer Time Zone

The time zone in which a resource is located.

Storage Virtualizer Location

The physical location of a storage virtualizer. The location is defined when a storage virtualizer is added to IBM Spectrum Control. You can add or edit the location of the storage virtualizer in the Properties pane for the storage system.

Storage Virtualizer Custom Tag 1, 2, and 3

User-defined text that is associated with a storage virtualizer. You can add or edit the custom tags for a storage virtualizer on the Properties page of the storage system.

Storage Virtualizer User Provided Name

The name that was specified for the storage virtualizer in the storage environment.

Storage Virtualizer Tier Type

The type of tier that is used on the storage virtualizer. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD). The tier is configured to be controlled by the Easy Tier® function.

This value is available only for resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Is Compression Active

Shows whether the compression feature is enabled on the storage virtualizer. If this value is **Yes**, the compression feature is enabled.

Storage Virtualizer Configuration

Shows that a storage virtualizer is configured to virtualize back-end storage. This value is always Storage Virtualizer.

Status information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Status

The condition of the resource, for example normal, warning, or error.

Storage Virtualizer Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Storage Virtualizer Last Data Collection

The date and time when storage statistics were last collected from the resource.

Storage Virtualizer Last Data Collection Status

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Type

The type of storage virtualizer. For example, the storage virtualizer can be an IBM Storwize V7000 system, an IBM SAN Volume Controller system, or another type of storage virtualizer.

Storage Virtualizer Vendor

The vendor who supplied the resource.

Storage Virtualizer Model

The model name or model number of the storage virtualizer.

Storage Virtualizer Machine Type

The machine type of a storage virtualizer. For example, the storage virtualizer can be Storwize V7000 - 2076, SAN Volume Controller - 2145, or another type of storage virtualizer.

Storage Virtualizer IP Address

The IP address of the resource.

Storage Virtualizer Serial Number

The serial number of the resource.

Data for storage virtualizer pools in capacity and relationship reports

You can include general information, capacity data, properties, and other information about storage virtualizer pools in capacity and relationship reports.

Information about storage virtualizer pools

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Pool Name

The name that was assigned to the pool when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Pool Capacity (GiB)

The total amount of storage space in a pool.

Storage Virtualizer Pool Available Space (GiB)

The amount of unused storage space in a storage virtualizer pool.

Storage Virtualizer Pool Total Volume Capacity (GiB)

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Virtualizer Pool Allocated Space (GiB)

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes.

The space that is allocated for thin-provisioned volumes is less than their virtual capacity. The Storage Virtualizer Pool Total Volume Capacity (GiB) property shows the virtual capacity of the volumes. If a pool does not contain thin-provisioned volumes, this value is the same as the value of the Storage Virtualizer Pool Total Volume Capacity (GiB) property.

This value is equal to Storage Virtualizer Used Space (GiB) for volumes that are not thin provisioned.

Storage Virtualizer Pool Used Volume Space (GiB)

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For resources that are running IBM Spectrum Virtualize, you can preallocate thin-provisioned volume space when volumes are created. For these resources, the Storage Virtualizer Pool Used Space might be different than the Storage Virtualizer Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Virtualizer Pool Used Space and Storage Virtualizer Pool Allocated Space are equal.

This value is accurate as of the most recent time that IBM Spectrum® Control collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Virtualizer Pool Extent Size (MiB)

The size of the extents that were specified when a pool was created. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Assigned Volume Space (GiB)

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Virtualizer Pool Unassigned Volume Space (GiB)

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Virtualizer Pool Tier Capacity (GiB)

The amount of pool tier storage of a particular type. The type can be hard disk drive (HDD) or solid-state drive (SSD) and is identified in the Storage Virtualizer Pool Tier Type property. The tier is configured to be controlled by the Easy Tier® function.

This value is available only for resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Pool Tier Available Space (GiB)

The amount of pool tier storage of a particular type that is unused. The type can be hard disk drive (HDD) or solid-state drive (SSD) and is identified in the Storage Virtualizer Pool Tier Type property. The tier is configured to be controlled by the Easy Tier function.

This value is available only for resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Pool Tier Number of MDisks

The number of managed disks that is on a storage virtualizer pool. The tier is configured to be controlled by the Easy Tier function.

This value is available only for resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Pool Number of Storage Volumes

The number of volumes in the storage virtualizer pool.

Storage Virtualizer Pool Number of MDisks

The number of managed disks on a resource.

Storage Virtualizer Pool Maximum I/O Capability

The projected maximum number of I/O operations per second for a pool. This value is calculated based on the value in the Storage Virtualizer Pool Back-End Storage Disks property, and on the values in following properties:

- Storage Virtualizer Pool Back-End Storage System Type
- Storage Virtualizer Pool Back-End Storage RAID Level
- Storage Virtualizer Pool Back-End Storage Disk Type

This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Pool Back-End Storage Disks

The number of physical disks that contribute to the volumes on the back-end storage system. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Pool Physical Allocation Percentage

The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Pool Virtual Allocation Percentage

The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. In this case, the virtual capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, use the Storage Virtualizer Pool Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

This value is only available for pools with thin-provisioned volumes.

Storage Virtualizer Pool Shortfall Percentage

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{total volume capacity} - \text{allocated space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Unallocated Volume Space (GiB)

The amount of the Total Volume Capacity in the pool that is not allocated.

IBM Spectrum Control uses the following formula to determine this percentage:

$$\text{total volume capacity} - \text{allocated space}$$

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool.

Storage Virtualizer Pool Unused Volume Space (GiB)

The amount of space that is allocated to the volumes in a pool and is not yet used.

IBM Spectrum Control uses the following formula to determine this value:

$$\text{allocated space} - \text{used pool space}$$

Storage Virtualizer Pool Unallocatable Volume Space (GiB)

The amount of space by which the Total Volume Capacity exceeds the physical capacity of a pool.

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool. If you create volumes with more virtual capacity than can be physically allocated in the pool, you can over commit storage in the pool.

This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

Storage Virtualizer Pool Used Space Percentage

The percentage of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

Storage Virtualizer Pool Compressed Virtual Capacity (GiB)

The total virtual capacity of all the volumes that are compressed in a storage virtualizer pool.

Storage Virtualizer Pool Compressed Capacity (GiB)

The amount of storage space that is used by compressed volumes in a storage virtualizer pool. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 20.

Storage Virtualizer Pool Uncompressed Capacity (GiB)

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Pool Is Encryptable

Shows whether the storage virtualizer pool can be encrypted. If this value is **Yes**, the storage virtualizer pool can be encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Is Encrypted

Shows whether the storage virtualizer pool is encrypted. If this value is **Yes**, the storage virtualizer pool is encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Is Solid State

Shows whether there are solid-state drives in the pool. This property can contain the following values:

Mixed

The pool contains both hard disk drives and solid-state drives.

Non solid state

The pool contains no solid-state drives.

Solid state

The pool contains a least one solid-state drive.

Storage Virtualizer Pool Is Thin Provisioned

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is **Yes**, the resource is thin-provisioned.

Storage Virtualizer Pool RAID Level

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Virtualizer Pool Warning Level

The percentage of used capacity of the resource at which a warning is generated. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Pool Tier

The tier level of pools on storage virtualizers. For example, the tier can be 1, 2, 3, or another tier number.

Storage Virtualizer Pool Tier Type

The type of tier that is used on the storage virtualizer pool. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD). The tier is configured to be controlled by the Easy Tier function.

This value is available only for resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Pool Is Easy Tier

Shows whether the IBM® System Storage® Easy Tier feature is enabled. The IBM System Storage Easy Tier feature can be set to **auto**, **on**, or **off**.

The IBM System Storage Easy Tier feature automatically helps optimize solid-state storage deployments in multitier systems.

Storage Virtualizer Pool Back-End Storage System Type

The type of storage system that is providing storage space to a pool. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Pool Back-End Storage RAID Level

The RAID level of the volumes on the back-end storage system that are providing storage space to a pool. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Pool Back-End Storage Disk Type

The class and speed of the physical disks that contribute to the volumes on the back-end storage system. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Pool Custom Tag 1, 2, and 3

User-defined text that is associated with a storage virtualizer pool. You can add or edit the custom tags for a storage virtualizer pool on the Properties notebook of the pool.

Storage Virtualizer Pool Is Compression Active

Shows whether the compression feature is enabled on the storage virtualizer pool. If this value is **Yes**, the compression feature is enabled.

Status information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Pool Status

The condition of the resource, for example normal, warning, or error.

Storage Virtualizer Pool Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Storage Virtualizer Pool Easy Tier Status

Shows whether the Easy Tier function is active or inactive. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Data for storage virtualizer volumes in capacity and relationship reports

You can include general information, capacity data, properties, information about copies of volumes, and other information about storage virtualizer volumes in capacity and relationship reports.

Information about storage virtualizer volumes

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Volume Name

The name that was assigned to the storage virtualizer volume when it was added to the system.

Storage Virtualizer Volume Assigned Host Connection

The host connection to which the storage virtualizer volume is assigned. The host connection is a definition in the storage virtualizer that contains the WWPN for the server. The storage virtualizer uses the WWPN to assign volumes to servers.

Storage Virtualizer Volume Host Connection OS Type

The operating system type of the server or hypervisor that the virtualizer volume is assigned to.

Storage Virtualizer Volume Preferred Node Name

The number of the storage virtualizer node that is used for I/O operations for the storage virtualizer volume.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Volume Capacity (GiB)

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

Storage Virtualizer Volume Allocated Space (GiB)

The amount of space that is reserved for a volume.

The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in Storage Virtualizer Capacity.

This value is equal to Storage Virtualizer Volume Used Space for the following resources:

- Resources that are not running IBM Spectrum Virtualize
- Resources that are running IBM Spectrum Virtualize and are not thin-provisioned

Storage Virtualizer Volume Used Space (GiB)

The amount of allocated space that is used by a volume.

For resources that are running IBM Spectrum Virtualize, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Virtualizer Volume Used Space (GiB) might be different than the Storage Virtualizer Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Virtualizer Volume Used Space (GiB) and Storage Virtualizer Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that IBM Spectrum® Control collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Virtualizer Volume Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Virtualizer Volume Used Allocated Space Percentage

The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{volume used space} \div \text{volume allocated space}) \times 100$$

This property is available only for volumes on resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Volume Unallocated Space (GiB)

The amount of space in a pool that is not reserved for a volume.

IBM Spectrum Control uses the following formula to determine the value:

capacity - allocated space

Storage Virtualizer Volume Physical Allocation Percentage

The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

IBM Spectrum Control uses the following formula to determine this percentage:

(allocated space ÷ capacity) × 100

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore the space that is reserved for volumes is 50 GiB.

Storage Virtualizer Volume Unused Space (GiB)

The amount of space that is allocated to a volume and is not yet used.

IBM Spectrum Control uses the following formula to determine this value:

allocated space - used space

This property is available only for pools on resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Volume Shortfall Percentage

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume. The higher the percentage, the more critical the shortfall of space.

IBM Spectrum Control uses the following formula to determine this percentage:

(unallocatable volume space ÷ unallocated volume space) × 100

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

(100 GiB unallocated volume space - 70 GiB remaining pool space) ÷ 100 GiB unallocated volume space × 100

This value is only available for pools with thin-provisioned volumes.

Storage Virtualizer Volume Unallocatable Space (GiB)

The unallocatable storage space in GiB of a thin-provisioned volume. Unallocatable space cannot be supplied by the pool to which the volume belongs.

Storage Virtualizer Volume Uncompressed Capacity (GiB)

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Storage Virtualizer Volume Tier Capacity (GiB)

The amount of tier storage of a particular type on a virtualizer volume. The type is identified in the Storage Virtualizer Volume Tier Type property. The tier is configured to be controlled by the Easy Tier® function.

Storage Virtualizer Volume Tier Available Space (GiB)

The amount of tier storage of a particular type that is unused on a virtualizer volume. The type is identified in the Storage Virtualizer Volume Tier Type property. The tier is configured to be controlled by the Easy Tier function.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Volume Type

The type of storage virtualizer volume. For example, the storage virtualizer volume can be striped, sequential, or another type of storage virtualizer volume. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume WWN

The worldwide name of the volume.

Storage Virtualizer Volume Fast Write State

Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted.

Storage Virtualizer Volume Grain Size (KiB)

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Volume Is Assigned

Shows whether the volume is assigned to a server. If this value is **Yes**, the volume is assigned to a server.

Storage Virtualizer Volume Is Auto Expand Enabled

Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is **Yes**, the volume or volume copy can allocate new extents from a pool automatically.

Storage Virtualizer Volume Is Encryptable

Shows whether the resource can be encrypted. If this value is **Yes**, the resource can be encrypted.

Storage Virtualizer Volume Is Encrypted

Shows whether the resource is encrypted. If this value is **Yes**, the resource is encrypted.

Storage Virtualizer Volume Is Formatted

Shows whether a storage virtualizer volume is formatted. If this value is **Yes**, the storage virtualizer volume is formatted.

Storage Virtualizer Volume Is Thin Provisioned

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is **Yes**, the resource is thin-provisioned.

Storage Virtualizer Volume Metro Mirror Name

The name of the Metro Mirror that keeps the synchronous copy of the storage virtualizer volume.

Storage Virtualizer Volume Mirror Count

- The number of mirrors that keep a synchronous copy of the resource.
- Storage Virtualizer Volume Throttle
 - The maximum number of commands that the storage virtualizer volume can queue.
- Storage Virtualizer Volume Warning Level
 - The percentage of used capacity of the resource at which a warning is generated.
- Storage Virtualizer Volume FlashCopy® Relationship
 - Shows whether a volume on a storage virtualizer is in a FlashCopy relationship. This property can contain the following values:
 - Source Relationship
 - The volume is the source of the relationship.
 - Target Relationship
 - The volume is the target of the relationship.
 - Both Source and Target Relationship
 - The volume is both the source and target of the relationship.
 - Not in a FlashCopy Relationship
 - The volume is not in a FlashCopy relationship.
- Storage Virtualizer Volume Remote Copy Relationship
 - Shows whether a volume on a storage virtualizer is in a remote copy relationship. This property can contain the following values:
 - Not in a Copy Relationship
 - The volume is not in a remote copy relationship.
 - Source Relationship
 - The volume is the source of the relationship.
 - Target Relationship
 - The volume is the target of the relationship.
 - Both Source and Target Relationship
 - The volume is both the source and target of the relationship.
- Storage Virtualizer Volume RAID Level
 - The RAID level of a storage virtualizer volume, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.
- Storage Virtualizer Volume Is Compressed
 - Shows whether the storage virtualizer volume is compressed. If this value is **Yes**, the volume is compressed.
- Storage Virtualizer Volume Is Easy Tier
 - Shows how the Easy Tier function is enabled on a virtualizer volume. For example, if the value of this column is auto, the Easy Tier function is enabled in automatic mode.
- Storage Virtualizer Volume Tier Type
 - The type of tier that is used on the virtualizer volume. The tier is configured to be controlled by the Easy Tier function.

Status information

You can create capacity and relationship reports that include the following information:

- Storage Virtualizer Volume Status
 - The condition of the resource, for example normal, warning, or error.
- Storage Virtualizer Volume Is Detected
 - Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.
- Storage Virtualizer Volume Native Status
 - Shows the level of access of nodes in the system to the storage virtualizer managed disk or volume. The level of access can be online, offline, degraded, or excluded.
- Storage Virtualizer Volume Easy Tier Status
 - Shows whether the Easy Tier function is active or inactive on the virtualizer volume.

Volume copy information

You can create capacity and relationship reports that include the following information:

- Storage Virtualizer Volume Copy ID
 - The identifier of a volume copy. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.
- Storage Virtualizer Volume Copy Grain Size (KiB)
 - The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.
- Storage Virtualizer Volume Copy Is Auto Expand Enabled
 - Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is **Yes**, the volume or volume copy can allocate new extents from a pool automatically.
- Storage Virtualizer Volume Copy Is Detected
 - Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.
- Storage Virtualizer Volume Copy Is Primary
 - Shows whether the volume copy is the primary volume copy. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.
- Storage Virtualizer Volume Copy Is Thin Provisioned
 - Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is **Yes**, the resource is thin-provisioned.
- Storage Virtualizer Volume Copy Total Number of Extents
 - The number of extents that are used by the storage virtualizer volume copy.
- Storage Virtualizer Volume Copy Type
 - The type of copy of the storage virtualizer volume. For example, the copy of the storage virtualizer volume can be sequential, a mirror copy, or another type of copy.
- Storage Virtualizer Volume Copy Warning Level
 - The percentage of used capacity of the resource at which a warning is generated.
- Storage Virtualizer Volume Copy Used Allocated Space Percentage

The percentage of reserved space for a volume copy that is being used. This value is always less than or equal to 100% because a volume copy cannot use more space than is allocated.

IBM Spectrum Control uses the following formula to determine this value:

$$(\text{used space} \div \text{allocated space}) \times 100$$

This property is available only for pools on resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Volume Copy Used Space (GiB)

The amount of allocated space that is used by a volume copy.

Storage Virtualizer Volume Copy Allocated Space (GiB)

The amount of space that is reserved for a volume copy.

The space that is allocated for a thin-provisioned volume copy is less than its virtual capacity, which is shown in Storage Virtualizer Capacity (GiB).

This value is equal to Storage Virtualizer Volume Copy Used Space (GiB) for volume copies that are not thin provisioned.

Storage Virtualizer Volume Copy Unallocated Space (GiB)

The amount of space in a pool that is not reserved for a volume copy.

IBM Spectrum Control uses the following formula to determine this value:

$$\text{capacity} - \text{allocated space}$$

Storage Virtualizer Volume Copy Physical Allocation Percentage

The percentage of physical space that is reserved for a volume copy. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

IBM Spectrum Control uses the following formula to determine this value:

$$(\text{allocated space} \div \text{capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volume copies is 50 GiB.

Storage Virtualizer Volume Copy Unused Space (GiB)

The amount of space that is allocated to a volume copy and is not yet used.

IBM Spectrum Control uses the following formula to determine this value:

$$\text{allocated space} - \text{used space}$$

This property is available only for pools on resources that are running IBM Spectrum Virtualize.

Storage Virtualizer Volume Copy Shortfall Percentage

The percentage of the remaining unallocated volume copy space in a pool that is not available to be allocated to a volume copy.

The higher the percentage, the more critical the shortfall of space.

IBM Spectrum Control uses the following formula to determine this value:

$$(\text{unallocatable space} \div (\text{capacity} - \text{real capacity})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

$$(\text{100 GiB unallocatable volume copy space} - \text{70 GiB remaining pool space}) \div \text{100 GiB unallocated volume space} \times 100$$

Available only for thin-provisioned volume copies.

Storage Virtualizer Volume Copy Unallocatable Space (GiB)

The unallocatable storage space in GiB of a thin-provisioned volume copy. Unallocatable space cannot be supplied by the pool to which the volume copy belongs.

Storage Virtualizer Volume Copy Is Compressed

Shows whether the volume copy is compressed. If this value is **Yes**, the volume copy is compressed.

Storage Virtualizer Volume Copy Uncompressed Capacity (GiB)

The amount of storage space that is used if the compressed space on the volume copy is uncompressed. For example, if 100 GiB of data is compressed, the value is 100.

Data for managed disks on storage virtualizers in capacity and relationship reports

You can include general information, capacity data, properties, and other information about managed disks on storage virtualizers in capacity and relationship reports.

Information about managed disks on storage virtualizers

You can create capacity and relationship reports that include the following information:

Storage Virtualizer MDisk Name

The name that was assigned to the managed disk on a storage virtualizer when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Virtualizer MDisk Capacity (GiB)

The amount of storage space on the managed disk.

Storage Virtualizer MDisk Available Space (GiB)

The amount of space on managed disks that is available in a pool, or is available to be added to a pool.

Storage Virtualizer MDisk Strip Size (KB)

The RAID strip size on a managed disk on a storage virtualizer. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer MDisk RAID Level

The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Spare Goal

The number of spare drives that are required to maintain redundancy. Use spare drives to protect against drive failures in the array on a managed disk on a storage virtualizer.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Type

The type of managed disk on a storage virtualizer. For example, the storage virtualizer can be a local or virtual managed disk.

Storage Virtualizer MDisk Mode

The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

Array

Extents are provided from local disks.

Managed

Extents are provided from a back-end storage volume.

Unmanaged

The managed disk is not used in the system.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Is Balanced

Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is **Yes**, the LUNs are balanced.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Fast Write State

Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer MDisk Write Verify

Shows whether all write operations on a managed disk on a storage virtualizer are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is **Yes**, all write operations are verified by a follow-up read operation.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Status information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer MDisk Status

The condition of the resource, for example normal, warning, or error.

Storage Virtualizer MDisk Native Status

Shows the level of access of nodes in the system to a managed disk or volume on a storage virtualizer. The level of access can be online, offline, degraded, or excluded.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Data for local disks on storage virtualizers in capacity and relationship reports

You can include general information, capacity data, properties, and other information about local disks on storage virtualizers in capacity and relationship reports.

Information about local disks on storage virtualizers

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Local Disk Name

The name that was assigned to the local disk on the storage virtualizer when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Local Disk Capacity (GiB)

The amount of storage space on the local disk.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Local Disk Is Encryptable

Shows whether the local disk can be encrypted. If this value is **Yes**, the local disk can be encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Is Encrypted

Shows whether the local disk is encrypted. If this value is **Yes**, the local disk is encrypted.

This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Storage Class

The storage technology of the local disk on a storage virtualizer. For example, the storage class can be a serial-attached SCSI (SAS) or solid-state drive (SSD). This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Is Solid State

Shows whether the disk is a solid-state drive. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Speed (RPM)

The speed of the local disk on a storage virtualizer. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Tag

A number that identifies the array to which a local disk belongs.

Status information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Local Disk Status

The condition of the resource, for example normal, warning, or error.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Local Disk Firmware Version

The version number of the firmware that is running on the local disk. This value is not shown for Hitachi Universal Storage Platform V storage systems that are configured to virtualize back-end storage.

Storage Virtualizer Local Disk Model

The model name or model number of the resource.

Storage Virtualizer Local Disk Serial Number

The serial number of the resource.

Storage Virtualizer Local Disk Vendor

The vendor who supplied the resource.

Data for storage virtualizer groups in capacity and relationship reports

You can include the name of storage virtualizer groups and properties for storage virtualizer groups in capacity and relationship reports.

Information about storage virtualizer groups

Storage virtualizer groups are deprecated in IBM Spectrum® Control 5.2.10 and later. Use storage resource groups instead. In the IBM Spectrum Control GUI, storage resource groups are available as general groups.

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Group Name

A user-defined name of the storage virtualizer group.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Virtualizer Group Creator

The user name of the user who created the storage virtualizer group.

Storage Virtualizer Group Last Modified

The date and time that the storage virtualizer group was last modified.

Storage Virtualizer Group Last Modified by User

The user name of the user who last modified the storage virtualizer group.

Data for storage systems in capacity and relationship reports

You can include general information, capacity data, properties of components, and other information about storage systems in capacity and relationship reports.

Information about storage systems

You can create capacity and relationship reports that include the following information:

Storage System Name

A user-defined name of the storage system. If a name was not defined, IBM Spectrum® Control shows the name that was defined when the storage system was added for monitoring.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage System Pool Capacity (GiB)

The amount of storage space in pools that are on the storage system. For an XIV® or IBM Spectrum Accelerate, this value represents the physical capacity of the pool, not the virtual capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage System Used Pool Space (GiB)

The amount of space that is in use in all pools on a storage system.

Storage System Pool Available Space (GiB)

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

Storage System Volume Capacity (GiB)

The amount of space on all volumes on the storage system.

Storage System Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage System Assigned Volume Space (GiB)

The amount of space on the storage system that is on volumes that are assigned to a server or storage virtualizer.

Storage System Unallocated Disk Space (GiB)

The amount of disk space that can be added to a pool.

Storage System Number of Disks

The number of disks on the storage system.

Storage System Number of Volumes

The number of volumes on a resource.

Storage System Cache (GiB)

The size of the cache on the storage system. This value is not shown for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage System Real Available Pool Space (GiB)

The amount of unused space in pools that have an associated soft size on an XIV or IBM Spectrum Accelerate. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Real Configured Pool Space (GiB)

The amount of storage space that is in pools that have an associated soft size on an XIV or IBM Spectrum Accelerate. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage System Volume Capacity for z/OS (GiB)

The amount of space on all volumes on the storage system that the z/OS® operating system can use.

Storage System Volume Capacity Assigned to MDisks (GiB)

The amount of space that is on volumes that are assigned to a storage virtualizer to use as managed disks.

Storage System Disk Capacity (GiB)

The amount of storage space on physical disks in a storage system, excluding spare disks. This value is not shown for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage System Total Disk Capacity (GiB)

The amount of space on physical disks on a storage system, including spare disks.

Storage System Physical Allocation Percentage

The percentage of physical space in storage system pools that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in the pools.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{allocated pool space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a 200 GiB storage pool. Therefore, the space that is reserved for volumes is 50 GiB.

Storage System Virtual Allocation Percentage

The percentage of physical space in storage system pools that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, for a total pool size of 15 GiB, the allocation percentage might be 200%. Therefore, the virtual capacity that is committed to the volumes in the pools is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pools. If the allocation percentage is 100% for the same pools, then the virtual capacity that is committed to the pools is 15 GiB. This configuration means that all the physical capacity of the pools is already allocated to volumes.

An allocation percentage that is higher than 100% is considered aggressive. The pools have insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pools. In such cases, use the Storage System Shortfall Percentage property to estimate how critical the shortage of space is for storage system pools.

Storage System Shortfall Percentage

The percentage of the remaining unallocated volume space in storage system pools that is not available to be allocated. The higher the percentage, the more critical the shortfall of pool space.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{volume space} - \text{used volume space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This value represents the percentage of the committed virtual space that is not available in pools. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of pools is 70 GiB, but 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that is unavailable is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage System MDisk Available Space (GiB)

The amount of storage space that is available on a managed disk. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage System Tier Available Space (GiB)

The amount of unused space in a tier. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage System Tier Capacity (GiB)

The total amount of storage space in a tier. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Component properties

You can create capacity and relationship reports that include the following information:

Storage System Access Information

The URL from which you can access the storage system or storage virtualizer.

Storage System Code Level

The Shared Ethernet Adapter level of a DS8000® storage system. For other storage systems, this value is the firmware version.

Storage System Location

The physical location of a storage system. The location is defined when a storage system is added to IBM Spectrum Control. You can add or edit the location of the storage system in the Properties pane of the storage system.

Storage System Custom Tag 1, 2, and 3

User-defined text that is associated with a storage system. You can add or edit the custom tags for a storage system in the Properties pane of the storage system.

Storage System User Provided Name

The name that was specified for the storage system in the storage environment.

Storage System Tier Type

The type of tier that is used on a storage system. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD).

This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage System Configuration

Shows whether a storage system is configured as back-end storage or configured to virtualize back-end storage. If this value is **Storage System**, the storage system is configured as back-end storage. If this value is **Storage Virtualizer**, the storage system is configured to virtualize back-end storage.

Storage System Is Compression Active

Shows whether the compression feature is enabled on the storage system. If this value is **Yes**, the compression feature is enabled.

Status information

You can create capacity and relationship reports that include the following information:

Storage System Status

The condition of the resource, for example normal, warning, or error.

Storage System Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Storage System Last Data Collection

The date and time when storage statistics were last collected from the resource.

Storage System Last Data Collection Status

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Storage System Type

The type of storage system. For example, the storage system can be an IBM System Storage DS8800 system, an IBM System Storage DS8700 system, an IBM System Storage XIV system, or another type of storage system.

Storage System Vendor

The vendor who supplied the resource.

Storage System Model

The model name or model number of the storage system.

Storage System Machine Type

The machine type of a storage system. For example, the storage system can be DS8800, DS8700, XIV - 2812, or another type of storage system.

Storage System IP Address

The IP address of the resource.

Storage System Serial Number

The serial number of the resource.

Data for storage system pools in capacity and relationship reports

You can include general information, capacity data, properties of components, and other information about storage system pools in capacity and relationship reports.

Information about storage systems pools

You can create capacity and relationship reports that include the following information:

Storage Pool Name

The name that was assigned to the pool when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Pool Capacity (GiB)

The total amount of storage space in a pool.

Storage Pool Available Space (GiB)

The amount of unused space that is not reserved for volumes in pools that are on the storage system.

IBM Spectrum® Control uses the following formula to determine this value:

$$\text{pool capacity} - \text{used space}$$

For XIV® systems and IBM Spectrum Accelerate, this value represents the unallocated physical space in the pool, not the unallocated virtual space. For some storage systems, this value usually includes only the usable capacity, but might also include overhead space if the pool is unformatted.

Storage Pool Assigned Volume Space (GiB)

The amount of space in the pool that is on volumes that are assigned to a server or storage virtualizer.

Storage Pool Unassigned Volume Space (GiB)

The amount of volume space in the pool that is not assigned to a server or storage virtualizer.

Storage Pool Total Volume Capacity (GiB)

The total storage space on all the volumes in a pool, which includes thin-provisioned and standard volumes. For thin-provisioned volumes, this value includes virtual space.

Storage Pool Real Configured Space (GiB)

The amount of storage space that is in an XIV or IBM Spectrum Accelerate pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage Pool Used Volume Space (GiB)

The amount of allocated space that is used by the volumes in a pool, which includes thin-provisioned and standard volumes.

For resources that are running IBM Spectrum Virtualize, you can preallocate thin-provisioned volume space when volumes are created. For these resources, the Storage Pool Used Space might be different than the Storage Pool Allocated Space for pools that contain thin-provisioned volumes. In other cases, the values for Storage Pool Used Space and Storage Pool Allocated Space are equal.

This value is accurate as of the most recent time that IBM Spectrum Control collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Pool Real Available Space (GiB)

The amount of unused space that is in an XIV or IBM Spectrum Accelerate pool that has an associated soft size. The soft size of a pool is the virtual size of a thin-provisioned pool.

Storage Pool Repository Capacity (GiB)

The amount of space on all extents in the repository of a pool. This space can be used to allocate track space-efficient volumes. This attribute applies only to DS8000® storage systems.

Storage Pool Used Repository Space (GiB)

The amount of space on all extents in the repository of a pool that are allocated. This space can be used to allocate track space-efficient volumes. This attribute applies only to the DS8000 storage systems.

Storage Pool Physical Allocation Percentage

The percentage of physical space in a pool that is reserved for volumes. This value is always less than or equal to 100% because you cannot reserve more physical space than is available in a pool.

IBM Spectrum Control uses the following formula to determine the allocation percentage:

$$(\text{allocated space} \div \text{pool capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Pool Virtual Allocation Percentage

The percentage of physical space in a pool that is committed to the total virtual capacity of the volumes in the pool. In thin-provisioned environments, this percentage exceeds 100% if a pool is overcommitted (over-provisioned).

IBM Spectrum Control uses the following formula to determine the allocation percentage:

$$(\text{total volume capacity} \div \text{pool capacity}) \times 100$$

For example, the allocation percentage is 200% for a total pool size of 15 GiB. Therefore, the virtual capacity that is committed to the volumes in the pool is 30 GiB. This configuration means that twice as much space is committed than is physically contained in the pool. If the allocation percentage is 100% for the same pool, then the virtual capacity that is committed to the pool is 15 GiB. This configuration means that all the physical capacity of the pool is already allocated to volumes. An allocation percentage that is higher than 100% is considered aggressive. The pool has insufficient physical capacity to satisfy the maximum allocation for all the thin-provisioned volumes in the pool. In such cases, you can use the value for Storage Pool Shortfall Percentage to estimate how critical the shortage of space is for a pool.

This value is only available for pools with thin-provisioned volumes.

Storage Pool Shortfall Percentage

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated.

The higher the percentage, the more critical the shortfall of pool space.

IBM Spectrum Control uses the following formula to determine this percentage:

$$(\text{unallocatable space} \div (\text{virtual capacity} - \text{allocated space})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by volumes while the pool capacity remains the same, this percentage increases.

For example, the physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to thin-provisioned volumes. If the volumes are using 50 GiB, then there is still 100 GiB committed to those volumes (150 GiB - 50 GiB). There is only 20 GiB of available pool space (70 GiB - 50 GiB). Because only 20 GiB of pool space is available, 80 GiB of the committed space cannot be allocated (100 GiB - 20 GiB). In this case, the percentage of committed space that cannot be allocated is 80% (80 GiB ÷ 100 GiB × 100).

This value is only available for pools with thin-provisioned volumes.

Storage Pool Allocated Space (GiB)

The amount of space that is reserved for all the volumes in a pool, which includes both thin-provisioned and standard volumes. The space that is allocated for thin-provisioned volumes is less than their virtual capacity, which is shown in the Storage Pool Total Volume Capacity (GiB) property. If a pool does not contain thin-

provisioned volumes, this value is the same as the value in the Storage Pool Total Volume Capacity (GiB) property. This value is equal to the value in the Storage Pool Used Volume Space (GiB) property for the following resources:

- Resources other than those that are running IBM Spectrum Virtualize and are configured as back-end storage
- Resources that are running IBM Spectrum Virtualize, are configured as back-end storage, and are not thin-provisioned

Storage Pool Extent Size (MiB)

The size of the extent that was specified when a pool was created. Smaller extent sizes limit the maximum size of the volumes that can be created in a pool. Smaller extent sizes minimize the amount of potentially wasted space per volume.

This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Pool Tier Capacity (GiB)

The information that this property shows depends on whether the report includes the Storage Pool Tier Type property:

- Report includes Storage Pool Tier Type: Shows the amount of space on solid-state drives and hard disk drives in a pool. The report shows the storage space on solid-state drives and on hard disk drives on different rows.
- Report does not include Storage Pool Tier Type: Shows the total storage space on all of the solid-state drives and hard disk drives in a pool.

The tier is configured to be controlled by the Easy Tier® function.

This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Pool Tier Available Space (GiB)

The information that this property shows depends on whether the report includes the Storage Pool Tier Type property:

- Report includes Storage Pool Tier Type: Shows the amount of unused space on solid-state drives and hard disk drives in a pool. The report shows the unused space on solid-state drives and on hard disk drives on different rows.
- Report does not include Storage Pool Tier Type: Shows the total unused space on all of the solid-state drives and hard disk drives in a pool.

The tier is configured to be controlled by the Easy Tier function.

This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Pool Tier Number of MDisks

The information that this property shows depends on whether the report includes the Storage Pool Tier Type property:

- Report includes Storage Pool Tier Type: Shows the number of managed disks on the tier that are solid-state drives or hard disk drives. The report shows the number of managed disks that are solid-state drives and the number that are hard disk drives on different rows.
- Report does not include Storage Pool Tier Type: Shows the total number of managed disks on all of the solid-state drives and hard disk drives in a tier.

This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Pool Unallocated Volume Space (GiB)

The amount of the Total Volume Capacity in the pool that is not allocated.

IBM Spectrum Control uses the following formula to determine this value:

total volume capacity - allocated space

The space that is allocated for thin-provisioned volumes is typically less than their virtual capacity. Therefore, the unallocated space represents the difference between the virtual capacity and the allocated space for all the volumes in the pool.

Storage Pool Unallocatable Volume Space (GiB)

The amount of space by which the Total Volume Capacity exceeds the physical capacity of a pool.

In thin-provisioned environments, it is possible to over commit (over provision) storage in a pool. If you create volumes with more virtual capacity than can be physically allocated in the pool, you can over commit storage in the pool.

This value represents the amount of volume space that cannot be allocated based on the current capacity of the pool.

Storage Pool Number of MDisks

The number of managed disks in a storage pool.

Storage Pool Compressed Virtual Capacity (GiB)

The total virtual capacity of all the volumes that are compressed in a pool.

Storage Pool Compressed Capacity (GiB)

The amount of storage space that is used by compressed volumes in a pool. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 20.

Storage Pool Uncompressed Capacity (GiB)

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Storage Pool Maximum I/O Capability

The projected maximum number of I/O operations per second for a pool. This value is calculated based on the value in the Storage Pool Back-End Storage Disks property, and on the values in following properties:

- Storage Pool Back-End Storage System Type
- Storage Pool Back-End Storage RAID Level
- Storage Pool Back-End Storage Disk Type

This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage Pool Back-End Storage Disks

The number of physical disks that contribute to the volumes on the back-end storage system. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Pool RAID Level

The RAID level of the resource, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Pool Is Solid State

Shows whether there are solid-state drives in the pool. This property can contain the following values:

Mixed

- The pool contains both hard disk drives and solid-state drives.
- Non solid state
 - The pool contains no solid-state drives.
- Solid state
 - The pool contains a least one solid-state drive.

Storage Pool Is Encrypted

Shows whether the resource is encrypted. If this value is **Yes**, the resource is encrypted.

Storage Pool Format

Shows the format of the volume or pool. The format can be a Count Key Data (CKD) format, or a fixed block format.

Storage Pool Is Thin Provisioned

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is **Yes**, the resource is thin-provisioned.

Storage Pool Logical Subsystem

The logical subsystem (LSS) to which a volume or pool belongs.

Storage Pool Rank Group

The rank group to which a pool is assigned, that is, 0 or 1. This property applies only to the DS8000 storage systems.

Storage Pool Lock Behavior

Shows whether the pool is locked for write operations, or is disabled for both read and write operations when storage space is limited. This property applies only to XIV systems and IBM Spectrum Accelerate.

Storage Pool Tier

The tier level of pools on storage systems. For example, the tier can be 1, 2, 3, or another tier number.

Storage Pool Custom Tag 1, 2, and 3

User-defined text that is associated with a storage pool. You can add or edit the custom tags for a storage pool on the Properties notebook of the pool.

Storage Pool Is Compression Active

Shows whether the compression feature is enabled on the storage pool. If this value is **Yes**, the compression feature is enabled.

Storage Pool Back-End Storage System Type

The type of storage system that provides storage space to a pool. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage Pool Back-End Storage Disk Type

The class and speed of the physical disks that contribute to the volumes on the back-end storage system. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Storage Pool Back-End Storage RAID Level

The RAID level of the volumes on the back-end storage system that provide storage space to a pool. This property applies only to pools on resources that are running IBM Spectrum Virtualize.

Status information

You can create capacity and relationship reports that include the following information:

Storage Pool Status

The condition of the resource, for example normal, warning, or error.

Storage Pool Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Storage Pool Tier Type

The type of tier that is used on the storage pool. For example, the tier type can be hard disk drive (HDD) or solid-state drive (SSD). The tier is configured to be controlled by the Easy Tier function.

This value is available only for resources that are running IBM Spectrum Virtualize.

Data for storage system volumes in capacity and relationship reports

You can include general information, capacity data, properties, information about copies of volumes, and other information about storage system volumes in capacity and relationship reports.

Information about storage system volumes

You can create capacity and relationship reports that include the following information:

Storage Volume Name

The name that was assigned to the storage volume when it was added to the system.

Storage Volume Assigned Host Connection

The host connection to which the storage volume is assigned. The host connection is a definition in the storage system that contains the WWPN for the server. The storage system uses the WWPN to assign volumes to servers.

Storage Volume Host Connection OS Type

The operating system type of the server or hypervisor that the volume is assigned to.

Storage System Volume Group ID

The unique identifier for a collection of volumes on DS8000® storage systems. The storage system generates this identifier.

Storage System Volume Group Name

The user-defined name for a collection of volumes on DS8000 storage systems.

Storage System Volume Preferred Node Name

The name of the node that is used for I/O operations for the storage system volume.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Volume Capacity (GiB)

The total amount of storage space that is committed to a volume.

For thin-provisioned volumes, this value represents the virtual capacity of the volume.

For XIV® systems and IBM Spectrum Accelerate, this value represents the physical (hard) capacity of the volume, not the virtual (soft) capacity. For other storage systems, this value might also include overhead space if the pool is unformatted.

Storage Volume Overhead (GiB)

The amount of space that is used for system management. The amount of space also includes space that is reserved for redundancy.

Storage Volume Allocated Space (GiB)

The amount of space that is reserved for a volume. The space that is allocated for a thin-provisioned volume is less than its virtual capacity, which is shown in the Storage Volume Capacity (GiB) property. This value is equal to the value in the Storage Volume Used Space (GiB) property for the following resources:

- Resources other than those that are running IBM Spectrum Virtualize and are configured as back-end storage
- Resources that are running IBM Spectrum Virtualize, are configured as back-end storage, and are not thin-provisioned

Storage Volume Used Space (GiB)

The amount of allocated space that is used by a volume.

For resources that are running IBM Spectrum Virtualize, you can preallocate thin-provisioned volume space when the volumes are created. In these cases, the Storage Volume Used Space (GiB) might be different than the Storage Volume Allocated Space (GiB). For volumes that are not thin provisioned, the values for Storage Volume Used Space (GiB) and Storage Volume Allocated Space (GiB) are equal.

This value is accurate as of the most recent time that IBM Spectrum® Control collected data about a volume. The value in this property might not be 100% accurate for the current state of volumes. This inaccuracy might occur because data collection is run on a set schedule and the used space on volumes can change rapidly.

Storage Volume Used Allocated Space Percentage

The percentage of reserved space for a volume that is being used. This value is always less than or equal to 100% because a volume cannot use more space than is allocated.

IBM Spectrum Control uses the following formula to determine this value:

$$(\text{volume used space} \div \text{volume allocated space}) \times 100$$

This property is available only for volumes on resources that are running IBM Spectrum Virtualize.

Storage Volume Unallocated Space (GiB)

The amount of space in a pool that is not reserved for a volume.

IBM Spectrum Control uses the following formula to determine this value:

$$\text{capacity} - \text{allocated space}$$

Available only for thin provisioned volumes.

Storage Volume Physical Allocation Percentage

The percentage of physical space that is reserved for a volume. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

IBM Spectrum Control uses the following formula to determine this value:

$$(\text{allocated space} \div \text{capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a volume size of 200 GiB. Therefore, the space that is reserved for volumes is 50 GiB.

Storage Volume Unused Space (GiB)

The amount of space that is allocated to a volume and is not yet used.

IBM Spectrum Control uses the following formula to determine this value:

$$\text{allocated space} - \text{used space}$$

This property is available only for pools on resources that are running IBM Spectrum Virtualize.

Storage Volume Shortfall Percentage

The percentage of the remaining unallocated volume space in a pool that is not available to be allocated to a volume.

The higher the percentage, the more critical the shortfall of space.

IBM Spectrum Control uses the following formula to determine this value:

$$(\text{unallocatable volume space} \div \text{volume unallocated space}) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in pools reaches a critically high level. For example, the physical space in pools might be less than the committed virtual space. In this case, the pools do not have enough space to fulfill the commitment to virtual space for a volume.

This Storage Volume Shortfall Percentage represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume is using 50 GiB, then there is still 100 GiB committed to that volume (150 GiB - 50 GiB). There is a shortfall of 30 GiB (70 GiB remaining pool space - 100 GiB remaining commitment of volume space to the volume). The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

$$\frac{100 \text{ GiB unallocated volume space} - 70 \text{ GiB remaining pool space}}{100 \text{ GiB unallocated volume space}} \times 100$$

Storage Volume Unallocatable Space (GiB)

The unallocatable storage space in GiB of a thin-provisioned volume. Unallocatable space cannot be supplied by the pool to which the volume belongs.

Storage Volume Uncompressed Capacity (GiB)

The amount of storage space that is used if the compressed volume space is uncompressed. For example, if 100 GiB of uncompressed data is compressed, and the size of the compressed data is 20 GiB, the value is 100.

Storage Volume Tier Capacity (GiB)

The amount of volume tier storage of a particular type. The type is identified in the Storage Volume Tier Type property. The tier is configured to be controlled by the Easy Tier® function.

Storage Volume Tier Available Space (GiB)

The amount of volume tier storage of a particular type that is unused. The type is identified in the Storage Volume Tier Type property. The tier is configured to be controlled by the Easy Tier function.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Volume Type

The type of storage volume. For example, the storage volume can be striped or sequential. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume WWN

The worldwide name of the volume.

Storage Volume Fast Write State

Shows whether the cache for a volume on a disk that is managed by a storage virtualizer is empty, contains data, or is corrupted.

Storage Volume Grain Size (KiB)

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Is Assigned

Shows whether the volume is assigned to a server. If this value is **Yes**, the volume is assigned to a server.

Storage Volume Is Auto Expand Enabled

Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is **Yes**, the volume or volume copy can allocate new extents from a pool automatically.

Storage Volume Is Encryptable

Shows whether the resource can be encrypted. If this value is **Yes**, the resource can be encrypted.

Storage Volume Is Encrypted

Shows whether the resource is encrypted. If this value is **Yes**, the resource is encrypted.

Storage Volume Is Formatted

Shows whether a storage volume is formatted. If this value is **Yes**, the volume is formatted. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage Volume Is Thin Provisioned

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is **Yes**, the resource is thin-provisioned.

Storage Volume Metro Mirror Name

The name of the Metro Mirror that keeps the synchronous copy of the storage virtualizer volume.

Storage Volume Mirror Count

The number of mirrors that keep a synchronous copy of the resource.

Storage Volume Throttle

The maximum number of commands that the storage volume can queue.

Storage Volume Warning Level

The percentage of used capacity of the resource at which a warning is generated.

Storage Volume FlashCopy® Relationship

Shows whether a volume on a storage system is in a FlashCopy relationship. This property can contain the following values:

Source Relationship

The volume is the source of the relationship.

Target Relationship

The volume is the target of the relationship.

Both Source and Target Relationship

The volume is both the source and target of the relationship.

Not in a FlashCopy Relationship

The volume is not in a FlashCopy relationship.

Storage Volume Remote Copy Relationship

Shows whether a volume on a storage system is in a remote copy relationship. This property can contain the following values:

Not in a Copy Relationship

The volume is not in a remote copy relationship.

Source Relationship

The volume is the source of the relationship.

Target Relationship

The volume is the target of the relationship.

Both Source and Target Relationship

The volume is both the source and target of the relationship.

Storage Volume RAID Level

The RAID level of a storage volume, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the resource.

Storage Volume Is Compressed

Shows whether the volume is compressed. If this value is **Yes**, the volume is compressed.

Storage Volume Is Volume Control Manager Database

Shows whether a volume is used as the Volume Control Manager Database. If this value is **Yes**, the volume is used as the Volume Control Manager Database.

Storage Volume Format

Shows the format of the volume or pool. The format can be a Count Key Data (CKD) format, or a fixed block format.

Storage Volume Serial Number

The serial number or volume ID of the volume.

Storage Volume Number

The number of the volume that is assigned by the system.

Storage Volume Logical Subsystem

The logical subsystem (LSS) to which a volume or pool belongs.

Storage Volume Is Easy Tier

Shows how the Easy Tier function is enabled on a volume. For example, if the value of this column is auto, the Easy Tier function is enabled in automatic mode.

Storage Volume Tier Type

The type of tier that is used on the volume. The tier is configured to be controlled by the Easy Tier function.

Status information

You can create capacity and relationship reports that include the following information:

Storage Volume Status

The condition of the resource, for example normal, warning, or error.

Storage Volume Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Storage Volume Is Easy Tier

Shows how the Easy Tier function is enabled on a volume. For example, if the value of this column is auto, the Easy Tier function is enabled in automatic mode.

Volume copy information

You can create capacity and relationship reports that include the following information:

Storage Volume Copy ID

The identifier of a volume copy.

Storage Volume Copy Grain Size (KiB)

The grain size with which a thin-provisioned volume was created. This value is typically 32, 64, 128, or 256 KiB. Larger grain sizes maximize performance, whereas smaller grain sizes maximize space efficiency. Grain sizes also limit the maximum virtual space of a volume.

Storage Volume Copy Is Auto Expand Enabled

Shows whether a volume or volume copy can allocate new extents from a pool automatically. Volumes or volume copies might allocate new extents to expand the real capacity of the storage virtualizer volume or volume copy. If this value is **Yes**, the volume or volume copy can allocate new extents from a pool automatically.

Storage Volume Copy Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Storage Volume Copy Is Primary

Shows whether the volume copy is the primary volume copy.

Storage Volume Copy Is Thin Provisioned

Shows whether a pool, volume, or volume copy is thin-provisioned. If this value is **Yes**, the resource is thin-provisioned.

Storage Volume Copy Total Number of Extents

The number of extents that are used by the storage volume copy.

Storage Volume Copy Type

The type of copy of the storage volume. For example, the copy of the storage volume can be sequential, a mirror copy, or another type of copy.

Storage Volume Copy Warning Level

The percentage of used capacity of the resource at which a warning is generated.

Storage Volume Copy Used Allocated Space Percentage

The percentage of reserved space for a volume copy that is being used. This value is always less than or equal to 100% because a volume copy cannot use more space than is allocated.

IBM Spectrum Control uses the following formula to determine this value:

$$(\text{used space} \div \text{allocated space}) \times 100$$

This property is available only for pools on resources that are running IBM Spectrum Virtualize.

Storage Volume Copy Used Space (GiB)

The amount of allocated space that is used by a volume copy.

Storage Volume Copy Allocated Space (GiB)

The amount of space that is reserved for a volume copy.

The space that is allocated for a thin-provisioned volume copy is less than its virtual capacity, which is shown in the Storage Virtualizer Capacity (GiB) property.

For volume copies that are not thin provisioned, this value is equal to the value of the Storage Volume Copy Used Space (GiB) property.

Storage Volume Copy Unallocated Space (GiB)

The amount of space in a pool that is not reserved for a volume copy.

IBM Spectrum Control uses the following formula to determine this value:

$$\text{capacity} - \text{allocated space}$$

Storage Volume Copy Physical Allocation Percentage

The percentage of physical space that is reserved for a volume copy. This value is always less than or equal to 100% because you cannot reserve more physical space than is available.

IBM Spectrum Control uses the following formula to determine this value:

$$(\text{allocated space} \div \text{capacity}) \times 100$$

For example, the physical allocation percentage is 25% for a total pool size of 200 GiB. Therefore, the space that is reserved for volume copies is 50 GiB.

Storage Volume Copy Unused Space (GiB)

The amount of space that is allocated to a volume copy and that is not yet used.

IBM Spectrum Control uses the following formula to determine this value:

$$\text{allocated space} - \text{used space}$$

This property is available only for pools on resources that are running IBM Spectrum Virtualize.

Storage Volume Copy Shortfall Percentage

The percentage of the remaining unallocated volume copy space in a pool that is not available to be allocated to a volume copy.

This value represents the percentage of the committed virtual space that is not available in a pool. As more space is used over time by a volume while the pool capacity remains the same, this percentage increases.

The higher the percentage value is, the more critical the shortfall of space is.

IBM Spectrum Control uses the following formula to determine this value:

$$(\text{unallocatable space} \div (\text{capacity} - \text{real capacity})) \times 100$$

You can use this percentage to determine when the amount of overcommitted space in a pool reaches a critically high level for the volume. For example, the physical space in a pool might be less than the committed virtual space. In this case, the pool does not have enough space to fulfill the commitment to virtual space for a volume.

For example, the remaining physical capacity of a pool is 70 GiB, however, 150 GiB of virtual space is committed to a thin-provisioned volume. If the volume uses 50 GiB, 100 GiB is, nevertheless, committed to that volume (150 GiB minus 50 GiB). There is a shortfall of 30 GiB when you use the following calculation:

$$(100 \text{ GiB remaining commitment of volume space to the volume} - 70 \text{ GiB remaining pool space})$$

The volume is overcommitted by 30 GiB based on the available space in the pool. The shortfall is 30% when you use the following calculation:

```
(100 GiB unallocatable volume copy space - 70 GiB
remaining pool space) ÷ 100 GiB unallocated volume space × 100
```

This property is available only for volume copies that are thin provisioned.

Storage Volume Copy Unallocatable Space (GiB)

The unallocatable storage space in GiB of a thin-provisioned volume copy. Unallocatable space cannot be supplied by the pool to which the volume copy belongs.

Storage Volume Copy Is Compressed

Shows whether the volume copy is compressed. If this value is **Yes**, the volume copy is compressed.

Storage Volume Copy Uncompressed Capacity (GiB)

The amount of storage space that is used if the compressed space on the volume copy is uncompressed. For example, if 100 GiB of data is compressed, the value is 100.

Data for storage system disks in capacity and relationship reports

You can include properties, and other data about storage system disks in capacity and relationship reports.

Information about storage system disks

You can create capacity and relationship reports that include the following information:

Storage Disk Name

The name that was assigned to the storage disk when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage Disk Capacity (GiB)

The amount of storage space that is on a physical disk on the storage system.

Component properties

You can create capacity and relationship reports that include the following information:

Storage Disk Flags

Shows whether a device is a spare device. This property applies only to DS8000® storage systems.

Storage Disk Is Encryptable

Shows whether the resource can be encrypted. If this value is **Yes**, the resource can be encrypted.

Storage Disk Is Encrypted

Shows whether the resource is encrypted. If this value is **Yes**, the resource is encrypted.

Storage Disk Is Solid State

Shows whether the disk is a solid-state drive.

Storage Disk Speed (RPM)

The speed of the disk.

Storage Disk Tag

A number that identifies the array to which a disk belongs. This property applies only to DS8000 storage systems.

Storage Disk Class

The storage technology of the disk. This property applies only to DS8000 storage systems. This property can contain the following values:

FC

The disk is a Fibre Channel disk.

Nearline ATA

The disk is a nearline Advanced Technology Attachment disk.

Nearline SAS

The disk is a nearline serial-attached SCSI disk.

SAS

The disk is a serial-attached SCSI disk.

SATA

The disk is a Serial Advanced Technology Attachment disk.

SSD

The disk is a solid-state drive.

Unsupported

The disk is on a system that is not a DS8000 series system.

Status information

You can create capacity and relationship reports that include the following information:

Storage Disk Status

The condition of the resource, for example normal, warning, or error.

Storage Disk Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Storage Disk Firmware Version

The version number of the firmware that is running on the disk.

Storage Disk Model

The model name or model number of the resource.

Storage Disk Serial Number

The serial number of the resource.

Storage Disk Vendor

The vendor who supplied the resource.

Data for storage system groups in capacity and relationship reports

You can include general information about, and properties of, storage system groups in capacity and relationship reports.

Information about storage system groups

Storage system groups are deprecated in IBM Spectrum® Control 5.2.10 and later. Use storage resource groups instead. In the IBM Spectrum Control GUI, storage resource groups are available as general groups.

You can create capacity and relationship reports that include the following information:

Storage System Group Name

The name that was assigned to the storage system group when it was added to the system.

Component properties

You can create capacity and relationship reports that include the following information:

Storage System Group Creator

The user name of the user who created the group.

Storage System Group Last Modified

The date and time that the resource was last modified.

Storage System Group Last Modified by User

The user name of the user who last modified the resource.

Data for managed disks on storage systems in capacity and relationship reports

You can include general information, capacity data, properties, and other information about managed disks on storage systems in capacity and relationship reports.

Information about managed disks on storage systems

You can create capacity and relationship reports that include the following information:

Storage System MDisk Name

The name that was assigned to the managed disk on a storage system when it was added to the system.

Capacity and usage data

You can create capacity and relationship reports that include the following information:

Storage System MDisk Capacity (GiB)

The amount of storage space on the managed disk.

Storage System MDisk Available Space (GiB)

The amount of storage space that is available on a managed disk. This value is only available for resources that are running IBM Spectrum Virtualize and are configured as back-end storage.

Storage System MDisk Strip Size (KB)

The RAID strip size on a managed disk on a storage system.

Component properties

You can create capacity and relationship reports that include the following information:

Storage MDisk RAID Level

The RAID level of the managed disk, such as RAID 5 or RAID 10. The RAID level affects the performance and fault tolerance of the volumes that are allocated from the managed disk.

Storage MDisk Spare Goal

The number of spare drives that are required to maintain redundancy. Use spare drives to protect the system against drive failures in the array on a managed disk on a storage virtualizer.

Storage System MDisk Type

The type of managed disk on a storage system. For example, the disk on a storage system can be a local managed disk.

Storage System MDisk Mode

The access mode of a managed disk on a storage virtualizer. The access mode describes how extents are provided for virtual disks. Extents can be provided to virtual disks in the following ways:

Array

Extents are provided from local disks.

Managed

Extents are provided from a back-end storage volume.

Unmanaged

The managed disk is not used in the system.

Storage System MDisk Is Balanced

Shows whether LUNs are balanced across storage controllers on the managed disk. If this value is **Yes**, the LUNs are balanced.

Storage System MDisk Fast Write State

Shows whether the cache for a volume on a disk that is managed by a storage system is empty, contains data, or is corrupted.

Storage System MDisk Write Verify

Shows whether all write operations on a managed disk on a storage system are verified by an immediate follow-up read operation. The follow-up read operation verifies that the write operation was successful. If this value is **Yes**, all write operations are verified by a follow-up read operation.

Status information

You can create capacity and relationship reports that include the following information:

Storage System MDisk Status

The condition of the resource, for example normal, warning, or error.

Storage MDisk Native Status

Shows the level of access of nodes in the system to a managed disk or volume on a storage system. The level of access can be online, offline, degraded, or excluded.

Data for storage resource groups in capacity and relationship reports

You can include general information about storage resource groups in capacity and relationship reports. In the IBM Spectrum® Control GUI, storage resource groups are available as general groups.

Information about storage resource groups

You can create capacity and relationship reports that include the following information:

Storage Resource Group Name

The name that was assigned to the storage resource group when it was created.

Storage Resource Group Description

A user-defined description of the storage resource group. The description is created by the user when the storage resource group is created from the IBM Spectrum Control CLI.

Storage Resource Group Last Modified by User

The user name of the user who last modified the resource.

Storage Resource Group Custom Tag 1, 2, and 3

User-defined text that is associated with a storage resource group.

Data for groups in capacity and relationship reports

You can include general information about groups in capacity and relationship reports. A group is a set of logically related volumes, file systems, and shares. For example, a group that represents a business critical application might include the volumes, file systems, and shares that provide storage to the application.

Information about groups

You can create capacity and relationship reports that include the following information:

Group Name

The name that was assigned to the group by the user when the group was created.

Group Description

A user-defined description of the group. The description is created by the user when the group is created from the IBM Spectrum® Control CLI.

Group Last Modified by User

The user name of the user who last modified the resource.

Group Custom Tag 1, 2, and 3

User-defined text that is associated with a group.

Group Type

User-defined text that specifies the type of group. For example, the group type might be Application, Department, Project, General, or a similar type.

Group Subtype

User-defined text that is associated with a group. For example, if the Group Type property is Application, the Group Subtype property might specify the type of application. If the Group Type property is DB2® Group, the Group Subtype property might be Logs, Tablespace, or another group that is related to DB2.

Data for switches in capacity and relationship reports

You can include properties and other information about switches in capacity and relationship reports.

Information about switches

You can create capacity and relationship reports that include the following information:

Switch Name

The logical name of the switch where a Fibre Channel port is located, or the name that was defined when the switch was added for monitoring. If neither name is available, IBM Spectrum® Control uses the WWN of the switch.

Component properties

You can create capacity and relationship reports that include the following information:

Switch WWN

The World Wide Name (WWN) of the switch. A WWN is the unique 64-bit identifier for the switch.

Parent Switch Name

The name of the parent switch of a logical switch. If a user-defined name is not set for the parent switch, the WWN is shown. If the switch is not a logical switch, no value is displayed.

Parent Switch WWN

The worldwide name of the parent switch.

Switch Management Telnet Address

The Telnet address that you access to manage the switch.

Switch Management SNMP Address

The Simple Network Management Protocol (SNMP) address that you access to manage the switch.

Switch Management URL Address

The URL that you access to manage the switch.

Switch Domain

The domain ID of a switch. The ID is an 8-bit identifier with a range of 0-255. This column is blank for physical switches that are parents of virtual switches.

Switch Firmware Version

The firmware version of the microcode on a switch.

Switch Location

The physical location of the switch. The location is defined when a switch is added to IBM Spectrum Control. You can add or edit the location of the switch on the General tab of the properties notebook.

Switch Custom Tag 1, 2, and 3

User-defined text that is associated with the switch. You can add or edit the custom tags for the switch on the General tab of the properties notebook.

Switch Data Source Count

The number of data sources that are managing a switch. A switch can be managed by multiple data sources, and a single data source can manage multiple switches. Data sources might include CIM agents, SNMP agents, or Storage Resource agents. IBM Spectrum Control communicates with a data source to collect information about a switch.

Only the data sources that were added to IBM Spectrum Control are included in this number. For example, a switch is managed by two CIM agents and one Storage Resource agent. If one of the CIM agents was not added to IBM Spectrum Control, then the value for this property is 2.

Switch Is Virtual

Shows whether a switch is a logical switch. If this value is **Yes**, the switch is a logical switch.

Switch Ports

The total number of ports on the switch that are not online, and ports that are online and connected to other ports.

Switch Connected Ports

The number of ports that are connected to a storage resource, where the storage resource can be a storage system, server, or hypervisor. This value is only available when a switch is viewed as a related resource of a storage system, server, or hypervisor.

Status information

You can create capacity and relationship reports that include the following information:

Switch Status

The condition of the resource, for example normal, warning, or error.

Switch Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Switch Last Data Collection

The date and time when storage statistics were last collected from the resource.

Switch Last Data Collection Status

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

Vendor, model, and device information

You can create capacity and relationship reports that include the following information:

Switch IP Address

The IP address of the resource.

Switch Serial Number

The serial number of the resource.

Switch Vendor

The vendor who supplied the resource.

Switch Model

The model name or model number of the resource.

Data for switch ports in capacity and relationship reports

You can include properties and other information about switch ports in capacity and relationship reports.

Information about switch ports

You can create capacity and relationship reports that include the following information:

Switch Name

The logical name of the switch where a Fibre Channel port is located, or the name that was defined when the switch was added for monitoring. If neither name is available, IBM Spectrum® Control uses the WWN of the switch.

Component properties

You can create capacity and relationship reports that include the following information:

Switch Port Number

The port number on the resource.

Switch Port WWPN

The worldwide port name of the port on the switch.

Switch Port Type

The type of port on the storage system, storage virtualizer, or switch. For example, the port type can be N_Port, F_Port, or another type of port.

Switch Port Speed (GiB/s)

The speed of a port, which is measured in GiB per second.

Switch Blade Slot Number

The number of the slot on the switch to which the blade is attached. This property applies to ports on blades.

Switch Port Domain, Port

The domain ID of a switch, followed by the port index of a switch. In switches that have blades, the port index can differ from the port number.

Status information

You can create capacity and relationship reports that include the following information:

Switch Port Status

The condition of the resource, for example normal, warning, or error.

Switch Port Is Detected

Shows whether the resource was detected the last time that data was collected. If this value is **Yes**, the resource was detected.

Switch Port Enabled State

Shows whether a port is enabled, disabled, or is enabled but offline.

Switch Port Operational Status

The current operational state of the switch port. For example, this value can be **Error** or **OK**.

Data for fabrics in capacity and relationship reports

You can include properties and other information about fabrics in capacity and relationship reports.

Information about fabrics

You can create capacity and relationship reports that include the following information:

Fabric Name

The user-defined name for the fabric, or the name that was assigned to the fabric by the vendor.

Fabric WWN

The worldwide name of the fabric.

Component properties

You can create capacity and relationship reports that include the following information:

Fabric Type

The vendor of the switches in a fabric, such as IBM®, Brocade, or Cisco. If there are different switch vendors in the same fabric, the fabric type is displayed as **Mixed**. If this information is unknown, the property does not contain a value.

Principal Switch of Fabric

The name of the principal switch in the fabric at the time of the last data collection. The fabric can change its principal switch dynamically when conditions require it.

Important: Cisco physical switches are part of a SAN and not a fabric, therefore this property does not contain a value for those switches.

Fabric Switches

The number of physical and virtual switches in the fabric.

Fabric Switch Ports

The number of ports that are on all the switches in the fabric.

Fabric Connected Switch Ports

The number of switch ports that are online and connected to other ports in a fabric. The value is not shown when a fabric is viewed as a related resource of a storage system, server, or hypervisor.

Parent Fabric Name

The user-defined name for the parent fabric, or the name that was assigned to the parent fabric by the vendor.

Fabric Active Zone Set

The name of the zone set that is active. A zone set is a collection of zones within a single fabric. Only one zone set can be active for a fabric.

Fabric Location

The physical location of the fabric. The location is defined when a fabric is added to IBM Spectrum® Control. You can add or edit the location of the fabric on the General tab of the properties notebook.

Fabric Data Source Count

The number of data sources that are managing a fabric. A fabric can be managed by multiple data sources, and a single data source can manage multiple fabrics. Data sources can include SMI agents for Brocade switches or SNMP agents for Cisco switches. IBM Spectrum Control communicates with a data source to collect information about a fabric.

Only the data sources that were added to IBM Spectrum Control are included in this number. For example, a fabric is managed by two SMI agents on separate switches, but only one of the switches was added to IBM Spectrum Control. The value for the Fabric Data Source Count property is 1, because IBM Spectrum Control is configured to communicate with just the SMI agent of the switch that was added.

Fabric Is Virtual

Shows whether a fabric is a virtual fabric. If this value is **Yes**, the fabric is a virtual fabric.

Fabric Custom Tag 1, 2, and 3

User-defined text that is associated with a fabric. You can add or edit the custom tags for the fabric on the General tab of the properties notebook.

Status information

You can create capacity and relationship reports that include the following information:

Fabric Status

The condition of the resource, for example normal, warning, or error.

Fabric Last Data Collection

The date and time when storage statistics were last collected from the resource.

Fabric Last Data Collection Status

The condition of the last data collection. The status can show if the collection was a success, a failure, or if data was collected from the resource.

Troubleshooting Cognos Analytics reports

Find answers to questions about resolving issues with reports in the optional Cognos® Analytics reporting tool.

Problem

You can set output format options for reports either when you schedule a report or run a report. If you choose PDF format for your report and the report contains many columns, all of the columns might not fit on one page. For example, if a report contains 20 columns, then the row for a specific resource might be shown on more than one page: 10 columns on the first page and 10 columns on the second page.


Solution

To fit the maximum number of columns on one page, set the PDF options in one of the following ways:

When you schedule a report

1. Navigate to a report in the Cognos Analytics reporting tool.
2. Right-click the report, then click Properties.
3. Click the Schedule tab for the report.
4. Click the New icon.
5. Click Select in the PDF options area.
6. From the Orientation list, select Landscape.
7. From the Paper size list, select 11 × 17.

When you run a report

1. Navigate to a report in the Cognos Analytics reporting tool.
2. Right-click the report, then click Edit report.
3. Click the Run options icon , then click Show run options.
4. From the Paper size list, select 11 × 17.
5. From the Paper orientation list, select Landscape.

Problem

When you create a custom report, you cannot find online help that describes the columns that you can include in your reports.

Solution

To view the online help, follow these steps:

1. Go to the URL for your Cognos Analytics server. The URL is similar to this URL: `http://myhostname:9300/bi`
2. Click Team Content in the Welcome portal.
3. Click IBM Spectrum Control Predefined Reports.
4. Click any predefined report.
5. Click the Help link at the upper right of the page.
6. Go to the custom reports help content.

To view information about custom reports, see [Custom reports about performance](#) and [Custom reports about capacity and relationships](#).

Problem

If the IBM Spectrum® Control database (TPCDB) is not accessible, you are prompted to reenter your user ID and password although you previously entered the correct user ID and password.

Solution


Check that the server that hosts TPCDB is available. Check that the firewall settings on the server are not preventing you from accessing TPCDB.

Problem

If the password to access the IBM Spectrum Control database (TPCDB) is changed, you must also change the password in the Cognos Analytics reporting tool.

Solution

In the Cognos Analytics reporting tool, complete the following steps:

1. Go to the URL for your Cognos Analytics server. The URL is similar to this URL: **http://myhostname:9300/bi**
2. Click Manage in the Welcome portal, then click Administration console.
3. Click the Configuration tab.
4. Click the TPCDB data source
5. Click the TPCDB connection. The breadcrumb text is **Directory > Cognos > TPCDB > TPCDB**.
6. In the Actions column, click the Set properties icon , and then click the Signon tab.
7. Click Edit the signon.
8. Make your changes, and then click OK.

Problem

When you use particular combinations of properties from the Storage Volume section in the Performance package in a custom report, a **No data available** message might be displayed. This occurs when you use the following combinations of properties in a custom report:

- Server Name, Hypervisor Name, and any other property from the Storage Volume section
- Server Name, Hypervisor Name, Hypervisor Cluster Name, and any other property from the Storage Volume section
- Server Name, Hypervisor Cluster Name, and any other property from the Storage Volume section

Solution

Do not use these combinations of properties in custom reports.

Problem

When you add a filter to a property in Query Studio, not all of the resources are listed in the Filter (Pick values from a list) area. For example, when you add a filter to the Storage Volume Name property, not all of the volumes are listed.

This problem occurs because the maximum number of items in a list in Query Studio is 5000. If there are more than 5000 items in the list, only the first 5000 values are displayed, in alphabetical order.

Solution

Click Search for values in the Filter (Pick values from a list) area. Type a keyword, and then click Search to search for volumes on a resource. For example, type a% to find volume names that begin with 'a' or 'A'.

Problem

When you run a report, it takes a long time for the report output to display. It might take a long time to run a report for some combinations of the interval and reporting period. For example, the report might take a long time if you select the Sample interval or Hourly interval with one of the following periods:

- Last 30 days
- Last 90 days
- All

Solution

When you run a report, use shorter reporting periods with the Sample interval or Hourly interval.

Problem

When you run a report, the output that is displayed contains no data. If you select the Daily interval and a recent reporting period, then no data might be available. For example, if you select Last 24 hours or Today, the data for those reporting periods is unlikely to be collected yet.

Solution

When you run a report with a Daily interval, use a reporting period other than Last 24 hours or Today.

Problem

If you create a custom performance report in Query Studio, in particular circumstances, the report takes a long time to complete. This problem might occur in the following circumstances:

- The report is for more than 10,000 volumes.
- The interval for the report is Hourly, and the data for the report was collected for more than 15 days.
- The interval for the report is Sample, and the data for the report was collected for more than 7 days.

Solution

When this problem occurs, complete the following steps for the report in Query Studio:

1. Click Run Report in the menu pane.
2. Click Advanced Options.
3. Clear each of the following check boxes:
 - Automatically generate footer summaries for measures
 - Automatically summarize detail values, suppressing duplicates
 - Enable drill up and drill down in the report output
 - Enable drill through from a package in the report output
4. Click OK.

Problem

If you used the IBM Spectrum Control Capacity and Relationships package in 5.2, 5.2.1, 5.2.2, 5.2.3, or 5.2.4 to create custom reports, you might need to re-create those reports.

Solution

If your custom reports use any of the following properties, you must re-create the reports in the Capacity and Relationships package:

- NAS File Set Name
- NAS File Set Comment
- NAS File Set Path
- NAS File Set Used Space (GiB)
- NAS File Set Used Inodes

- NAS File Set Is Detected
- NAS File Set Status

To view information about how to create custom reports in the Capacity and Relationships package, see [Creating custom capacity and relationship reports](#).

Showing the package version number and build ID

Access information about the packages for custom reports and predefined reports, such as the version number and the build ID.

About this task

You can identify the version number and build ID for each of the following packages:

- Capacity and Relationships
- Performance

Procedure

1. Go to the URL for your Cognos® Analytics server. The format of the URL is similar to this URL: `http://myhostname:9300/bi`
2. Click Team Content in the Welcome portal.
3. Click IBM Spectrum Control Packages.
4. Depending on the package whose version number and build ID you want to see, complete one of the following steps:
 - To identify the version number and build ID of the Capacity and Relationships package, right-click Capacity and Relationships, then click Create report.
 - To identify the version number and build ID of the Performance package, right-click Performance, then click Create report.
5. In the navigation tree, position the cursor over the item under the name of the package.

Results

The version number and the build ID for the package is shown.

Reference

View reference information that is related to IBM Spectrum® Control. Topics include information about alerts, fabrics, commands, configuration and log files, performance metrics, protocols, standards, and accessibility features.

- [Return codes used by Storage Resource agent](#)
This topic lists the return codes used by the Storage Resource agent.
- [Agent types for monitoring fabrics and switches](#)
Depending on the type of switch you want to manage, you can use a CIM agent or SNMP agent as the data source for the switch.
- [Supported storage systems providing full disk encryption and solid-state drives](#)
IBM Spectrum Control supports full disk encryption and solid-state drives in the IBM® System Storage® DS8000® systems described in this topic.
- [Performance counters](#)
Use the metrics based on performance counters to measure and view the performance of storage systems and switches.
- [Command-line interface](#)
The following sections describe the IBM Spectrum Control command-line interface (CLI).
- [agent.sh command](#)
The `agent.sh` command lets you start, stop, and restart the Storage Resource agent. You can also display the status and version of the Storage Resource agent.
- [dataCollector command](#)
The `dataCollector` command is used to start and stop the data collector service.
- [Configuration files](#)
Use the parameters in IBM Spectrum Control configuration files to help resolve problems.
- [Log files](#)
When you have a problem, you can check several product log files.
- [Script parameters](#)
Script parameters provide specific information on the alert that triggered the script to be run.
- [Opening IBM Spectrum Control on Windows operating systems](#)
You can open IBM Spectrum Control CLIs and GUIs and administer IBM Spectrum Control on Windows operating systems.
- [Windows services used by IBM Spectrum Control](#)
To start, stop, or restart a component or related program in IBM Spectrum Control, use the Windows Services panel.
- [Frequently Asked Questions](#)
View answers to common questions about IBM Spectrum Control.
- [Protocols and standards](#)
This section provides an overview of the protocols and standards that are used within IBM Spectrum Control.
- [IBM Spectrum Control technical community](#)
Connect, learn, and share with storage professionals: product support technical experts who provide their perspectives and expertise.
- [Accessibility features for IBM Spectrum Control](#)
Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Return codes used by Storage Resource agent

This topic lists the return codes used by the Storage Resource agent.

The following table lists the return codes used by the Storage Resource agent during installation, uninstallation, and upgrade.

Table 1. Storage resource agent return codes

Return code	Explanation
1	There is a problem uninstalling the Fabric agent. For more information about what caused the uninstallation to fail, check the uninstallation logs for the Fabric agent.
2	Command not valid.
3	Option provided is not valid.
5	Argument is not valid.
6	Missing value for argument (e.g. -installLoc <Value>, where <Value> is missing).
7	Missing localized string in message file.
8	Probe is running.
9	Failed to open file for write.
10	Failed to close file.
11	Logfile not specified.
19	Failed in tracing.
21	Cannot spawn a probe because it is busy.
30	This is an internal error in initializing tracing. Save the error message and error log and contact your service representative for assistance in resolving the error.
32	Invalid socket.
33	Start service failed.
34	Registry entry not found.
35	Deployment of the agent failed, error creating startup scripts.
36	File does not exist.
40	Missing upgrade files.
41	Failed to extract files in upgrade process.
42	Failed to stop probe in upgrade process.
43	Failed to stop Agent.
44	Failed to start Agent.
45	Agent Registration to server failed.
46	File extraction needs more space.
47	Failed to open archive file.
48	Agent did not start after upgrade.
49	Installation directory not valid at upgrade time.
50	Probe is running.
51	Data file not found.
52	Exit code not in the output file.
53	Failed to send job status.
54	Failed to copy certificate files.
55	Failed to create directory.
56	Failed to remove directory.
57	Exec command failed.
58	Conversion of wide character failed.
59	Installation directory not valid.
60	Server name not defined.
61	Error in removing entries from configuration file.
62	Failed to stop probe at uninstall time.
63	Failed to remove registry entry at uninstall time.
64	Failed to remove service entry at uninstall time.
65	Failed to stop service at uninstall time.
66	Specified server name is not valid.
67	There is an error installing the Storage Resource agent as part of the migration process. For more information about what caused the installation to fail, check the installation logs for the Storage Resource agent.
71	Failed to spawn process.
73	Failed to spawn probe process.
101	Failed to create lock at installation time.
102	Failed to stop probe at re-installation time.
103	Failed to stop agent at re-installation time.
104	Failed to create registry at installation time.
105	Failed to extract files at installation time.
106	Failed to create entries in configuration file at installation time.
107	Failed to stop service at re-installation time.
108	Service already exists.
109	Failed to create service.
110	Failed to start service.
111	Probe failed at installation time.
112	Creation of daemon failed.
113	Installation of GUID failed.

Return code	Explanation
114	Commttype parameter is not valid.
115	Specified port is in use.
116	Installation/Upgrade does not have enough space.
117	Installation in progress.
118	Cannot get server name from Server.
119	Installation location is not empty.
120	Missing parameter Server Name.
121	Missing parameter Server Port.
123	Missing parameter Server IP.
124	Missing parameter Agent Port.
125	Missing parameter Installation location.
126	A value has not been specified for parameter userID. A value for this parameter is required when using RXA-based communication to deploy a Storage Resource agent as a non-daemon service. IBM Spectrum® Control uses this user ID when connecting to the computer on which the agent will be installed.
127	Deployment from Windows to Linux® failed.
130	Failed to send probe results.
131	Failed to initialize Agent.
133	Missing port number for Service.
134	Get data file stat failed.
135	Get data file read failed.
137	Failed to send data to server.
138	Failed to receive data from server.
139	Full path not specified for copy file.
140	Create file failed in copy file function.
141	Write file failed in copy file function.
142	Open file failed in copy file function.
143	Read file failed in copy file function.
145	UCS conversion failed.
146	Server connection failed.
148	Failed to create zip file.
149	Failed to unzip file.
160	Failed to send scan results.
161	Failed to send TSM status results.
164	Failed to validate user.
165	Job file was not found.
166	Job was not stopped.
168	Not enough free space available while copying file from server.
169	Multipath driver not found.
170	Multipath device not found.
171	Multipath policy is not supported.
172	Only Round Robin policy is supported for Multipath DM driver.
175	The command failed to run.
176	The command ran, but failed to complete successfully.

Agent types for monitoring fabrics and switches

Depending on the type of switch you want to manage, you can use a CIM agent or SNMP agent as the data source for the switch.

Table 1. Agent types for switch and fabric functions

Function	Brocade	Cisco
Monitor performance	REST API or SMI agent	SNMP agent
Collect information about switches and switch ports	REST API or SMI agent	SNMP agent
Collect information about topology connectivity	REST API or SMI agent	SNMP agent
Collect information about zoning information and zone control	REST API or SMI agent	SNMP agent
Generate alerts	REST API or SMI agent	SNMP agent

Check the [IBM Spectrum Control interoperability matrix for switches](#) for information about the switches and directors that are supported by IBM Spectrum® Control, and limitations that you need to know about when you use these devices.

Supported storage systems providing full disk encryption and solid-state drives

IBM Spectrum® Control supports full disk encryption and solid-state drives in the IBM® System Storage® DS8000® systems described in this topic.

IBM Spectrum Control supports full disk encryption and solid-state drives in DS8000 4.2 and later.

Performance counters

Use the metrics based on performance counters to measure and view the performance of storage systems and switches.

Overview

Performance counters are available for the following resources:

- [Volume counters](#)
- [Disk counters](#)
- [Storage system and switch port counters](#)

Volume counters

Volume performance counters are divided into the following categories:

- [Volume I/O counters](#)
- [Volume cache hit counters](#)
- [Volume data counters](#)
- [Volume service time counters](#)
- [Volume cache counters](#)
- [Volume miscellaneous counters](#)
- [Volume special computed values](#)

Table 1. Volume I/O counters

Counter	Counter type	Resource or component type	Description
Read I/O Count (normal)	001	DS8000®: volume, array, controller, storage system	Number of input/output (I/O) operations for nonsequential read requests, for a particular component over a particular time interval.
Read I/O Count (sequential)	002	DS8000: volume, array, controller, storage system	Number of I/O operations for sequential read requests, for a particular component over a particular time interval.
Read I/O Count (overall)	003	IBM Spectrum Virtualize: volume, node, I/O group, pool, storage system SMI-S BSP: volume, controller, storage system XIV®: volume, module, storage system	Number of I/O operations for all read requests, for a particular component over a particular time interval.
Write I/O Count (normal)	004	DS8000: volume, array, controller, storage system	Number of I/O operations for nonsequential write requests, for a particular component over a particular time interval.
Write I/O Count (sequential)	005	DS8000: volume, array, controller, storage system	Number of I/O operations for sequential write requests, for a particular component over a particular time interval.
Write I/O Count (overall)	006	IBM Spectrum Virtualize: volume, node, I/O group, pool, storage system SMI-S BSP: volume, controller, storage system XIV: volume, module, storage system	Number of I/O operations for all write requests, for a particular component over a particular time interval.
Total I/O Count (overall)	009	SMI-S BSP: volume, controller, storage system	Total number of I/O operations for all read and write requests, both sequential and nonsequential, for a particular component over a particular time interval.
Read Track Count	062	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of cache track read operations, for a particular component over a particular time interval.
Write Track Count	063	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of cache track write operations, for a particular component over a particular time interval.
Fast Write Mode Track Count	065	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of tracks writes that were processed in fast-write write mode, for a particular component over a particular time interval. Included as part of the "Write Track Count".
Write Through Mode Track Count	066	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of tracks writes that were processed in write-through write mode, for a particular component over a particular time interval. Included as part of the "Write Track Count".
Flush Through Mode Track Count	067	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of tracks writes that were processed in flush-through write mode, for a particular component over a particular time interval. Included as part of the "Write Track Count".
Cache Overflow Mode Track Count	068	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of tracks writes that were processed in cache overflow write mode, for a particular component over a particular time interval. Included as part of the "Write Track Count".
Global Mirror Write Count	090	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of write operations issued to the secondary site for Global Mirror, for a particular component over a particular time interval.

Counter	Counter type	Resource or component type	Description
Global Mirror Overlapping Write Count	091	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of serialized overlapping write operations encountered at the primary site for Global Mirror, for a particular component over a particular time interval.
HPF Read I/O Count	093	DS8000: volume, array, controller, storage system	Number of read I/O operations (both normal and sequential) that were issued using the High Performance FICON® (HPF) feature of the storage system, for a particular component over a particular time interval.
HPF Write I/O Count	094	DS8000: volume, array, controller, storage system	Number of write I/O operations (both normal and sequential) that were issued using the High Performance FICON (HPF) feature of the storage system, for a particular component over a particular time interval.
Non-Preferred Node I/O Count	096	IBM Spectrum Virtualize: I/O group, volume	Number of I/O operations (both read and write) that were performed using the non-preferred nodes of the volumes. This value is a subset of the total I/O operations for the component, and is normally expected to be 0.
Maintenance Operations Count	098	BSP: volume, controller, storage system	Cumulative count of all disk maintenance operations (SCSI commands such as: Verify, skip-mask, XOR read, XOR write-read, and so on). This value is required to understand the load on the component that might interfere with normal read and write operations.
Read Hit I/O Percentage for Small Transfers	104	XIV: volume, module, storage system	Percentage of I/O read operations over a specified interval that resulted in cache hits. Applies to data transfers that are ≤ 8 KB.
Read Hit I/O Percentage for Medium Transfers	105	XIV: volume, module, storage system	Percentage of I/O read operations over a specified interval that resulted in cache hits. Applies to data transfers that are > 8 KB and ≤ 64 KB.
Read Hit I/O Percentage for Large Transfers	106	XIV: volume, module, storage system	Percentage of I/O read operations over a specified interval that resulted in cache hits. Applies to data transfers that are > 64 KB and ≤ 512 KB.
Read Hit I/O Percentage for Huge Transfers	107	XIV: volume, module, storage system	Percentage of I/O read operations over a specified interval that resulted in cache hits. Applies to data transfers that are > 512 KB.
Read Miss I/O Percentage for Small Transfers	108	XIV: volume, module, storage system	Percentage of I/O read operations over a specified interval that resulted in cache misses. Applies to data transfers that are ≤ 8 KB.
Read Miss I/O Percentage for Medium Transfers	109	XIV: volume, module, storage system	Percentage of I/O read operations over a specified interval that resulted in cache misses. Applies to data transfers that are > 8 KB and ≤ 64 KB.
Read Miss I/O Percentage for Large Transfers	110	XIV: volume, module, storage system	Percentage of I/O read operations over a specified interval that resulted in cache misses. Applies to data transfers that are > 64 KB and ≤ 512 KB.
Read Miss I/O Percentage for Huge Transfers	111	XIV: volume, module, storage system	Percentage of I/O read operations over a specified interval that resulted in cache misses. Applies to data transfers that are > 512 KB.
Write Hit I/O Percentage for Small Transfers	112	XIV: volume, module, storage system	Percentage of I/O write operations over a specified interval that resulted in cache hits. Applies to data transfers that are ≤ 8 KB.
Write Hit I/O Percentage for Medium Transfers	113	XIV: volume, module, storage system	Percentage of I/O write operations over a specified interval that resulted in cache hits. Applies to data transfers that are > 8 KB and ≤ 64 KB.
Write Hit I/O Percentage for Large Transfers	114	XIV: volume, module, storage system	Percentage of I/O write operations over a specified interval that resulted in cache hits. Applies to data transfers that are > 64 KB and ≤ 512 KB.
Write Hit I/O Percentage for Huge Transfers	115	XIV: volume, module, storage system	Percentage of I/O write operations over a specified interval that resulted in cache hits. Applies to data transfers that are > 512 KB.
Write Miss I/O Percentage for Small Transfers	116	XIV: volume, module, storage system	Percentage of I/O write operations over a specified interval that resulted in cache misses. Applies to data transfers that are ≤ 8 KB.
Write Miss I/O Percentage for Medium Transfers	117	XIV: volume, module, storage system	Percentage of I/O write operations over a specified interval that resulted in cache misses. Applies to data transfers that are > 8 KB and ≤ 64 KB.
Write Miss I/O Percentage for Large Transfers	118	XIV: volume, module, storage system	Percentage of I/O write operations over a specified interval that resulted in cache misses. Applies to data transfers that are > 64 KB and ≤ 512 KB.
Write Miss I/O Percentage for Huge Transfers	119	XIV: volume, module, storage system	Percentage of I/O write operations over a specified interval that resulted in cache misses. Applies to data transfers that are > 512 KB.
Read SSD Hit Percentage for Small Transfer Sizes	153	XIV: volume, pool, module, host connection, storage system	Percentage of read operations over a specified time interval that accessed SSD cache memory. Applies to data transfers that are ≤ 8 KB.
Read SSD Hit Percentage for Medium Transfer Sizes	154	XIV: volume, pool, module, host connection, storage system	Percentage of read operations over a specified time interval that accessed SSD cache memory. Applies to data transfers that are > 8 KB and ≤ 64 KB.

Counter	Counter type	Resource or component type	Description
Read SSD Hit Percentage for Large Transfer Sizes	155	XIV: volume, pool, module, host connection, storage system	Percentage of read operations over a specified time interval that accessed SSD cache memory. Applies to data transfers that are > 64 KB and ≤ 512 KB.
Read SSD Hit Percentage for Huge Transfer Sizes	156	XIV: volume, pool, module, host connection, storage system	Percentage of read operations over a specified time interval that accessed SSD cache memory. Applies to data transfers that are > 512 KB.
Unaligned Write I/O Count	180	IBM Spectrum Virtualize ² : node, I/O group, host connection, storage system	The number of write operations that are not aligned on a boundary between 4 KB physical blocks. Unaligned write operations can cause a significant decrease in efficiency of write operations to some types of back-end disks. You can ignore this counter for iSeries servers.
Compressed Volumes I/O Count	182	IBM Spectrum Virtualize: node, I/O group, storage system	The total number of all read and write operations for compressed volumes.
Cache Read I/O Count - Volume Cache	610	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of read I/Os measured at the volume cache, which includes host reads as well as reads needed for remote replication.
Cache Read I/O Count - Volume Copy Cache	609	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of read I/Os measured at the volume copy cache, which includes host reads as well as reads needed for replication and other internal processing.
Cache Write I/O Count - Volume Cache	612	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of write I/Os measured at the volume cache, which includes host writes as well as writes needed for remote replication.
Cache Write I/O Count - Volume Copy Cache	611	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of write I/Os measured at the volume copy cache, which includes host writes as well as writes needed for replication and other internal processing.
Cache Prestage I/O Count - Volume Copy Cache	613	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of prefetch disk-to-cache transfer operations that are processed in the volume copy cache.
Cache Stage I/O Count - Volume Cache	615	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of disk-to-cache transfer operations that are processed in the volume cache.
Cache Stage I/O Count - Volume Copy Cache	614	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of disk-to-cache transfer operations that are not prefetch operations and are processed in the volume copy cache.
Cache Destage I/O Count - Volume Cache	617	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache-to-disk transfer operations that are processed in the volume cache.
Cache Destage I/O Count - Volume Copy Cache	616	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache-to-disk transfer operations that are processed in the volume copy cache.
In-flight I/O Count	658	IBM Spectrum Virtualize ¹ : managed disk	The actual number of concurrent operations that are issued from the cache to the managed disk. The number of in-flight operations cannot exceed the target, but also varies over time, depending on the type of I/O operations and on the performance of the managed disk.
Target I/O Count	657	IBM Spectrum Virtualize ¹ : managed disk	The maximum number of concurrent operations that can be issued from the cache to the managed disk. This target changes over time, depending on changes in cache utilization and on changes in the performance of the managed disk.
Notes:			
1. This counter is only available for resources that are running IBM Spectrum Virtualize 7.3 or later.			
2. This counter is only available for resources that are running IBM Spectrum Virtualize 7.4 or later.			

Table 2. Volume cache hit counters

Counter	Counter type	Resource or component type	Description
Read Cache Hits (normal)	010	DS8000: volume, array, controller, storage system	Number of cache hits for nonsequential read requests, for a particular component over a particular time interval.
Read Cache Hits (sequential)	011	DS8000: volume, array, controller, storage system	Number of cache hits for sequential read requests, for a particular component over a particular time interval.
Read Cache Hits (overall)	012	SMI-S BSP: volume, controller, storage system XIV: volume, module, storage system	Number of cache hits for all read requests, for a particular component over a particular time interval.
Write Cache Hits (normal)	013	DS8000: volume, array, controller, storage system	Number of cache hits for nonsequential write requests, for a particular component over a particular time interval.

Counter	Counter type	Resource or component type	Description
Write Cache Hits (sequential)	014	DS8000: volume, array, controller, storage system	Number of cache hits for sequential write requests, for a particular component over a particular time interval.
Write Cache Hits (overall)	015	SMI-S BSP: volume, controller, storage system XIV: volume, module, storage system	Number of cache hits for all write requests, for a particular component over a particular time interval.
Read Tracks Cache Hits	069	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of read track cache hits that were entirely serviced from prestaged or non-prestaged tracks for a particular component over a particular time interval.
Read Track Cache Hits Sector Count	070	IBM Spectrum Virtualize: node	Number of sectors read for cache hits that were entirely serviced from prestaged or non-prestaged tracks for a particular component over a particular time interval.
Read Sector Cache Hits - Volume Cache	69	IBM Spectrum Virtualize: volume, pool, node, I/O group, host connection, storage system	The number of cache sectors or tracks that achieved read cache hits in the volume cache. For resources that are running a version of IBM Spectrum Virtualize that is earlier than 7.3, the unit of the counter is tracks.
Read Sector Cache Hits - Volume Copy Cache	607	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache sectors that achieved read cache hits in the volume copy cache.
Prestage Cache Hits	80	IBM Spectrum Virtualize: volume, pool, node, I/O group, host connection, storage system	The number of prefetched cache sectors or tracks that achieved read cache hits in the volume copy cache. For resources that are running a version of IBM Spectrum Virtualize that is earlier than 7.3, the unit of the counter is tracks.
Destage Cache Hits - Volume Cache	81	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of destaged cache sectors or tracks that achieved write cache hits in the volume cache. That is, the number of writes that modified only data that was already marked "dirty" in the cache. For resources that are running a version of IBM Spectrum Virtualize that is earlier than 7.3, the unit of the counter is tracks.
Destage Cache Hits - Volume Copy Cache	608	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of destaged cache sectors that achieved write cache hits in the volume copy cache. That is, the number of writes that modified only data that was already marked "dirty" in the cache.
Note:			
1. This counter is only available for resources that are running IBM Spectrum Virtualize 7.3 or later.			

Table 3. Volume data counters

Counter	Counter type	Resource or component type	Description
Read KB Count	019	IBM Spectrum Virtualize: volume, node, I/O group, pool, storage system SMI-S BSP: volume, controller, storage system XIV: volume, module, storage system	Number of KB (2 ¹⁰ bytes) of data transferred for read operations, for a particular component over a particular time interval.
Write KB Count	020	IBM Spectrum Virtualize: volume, node, I/O group, pool, storage system SMI-S BSP: volume, controller, storage system XIV: volume, module, storage system	Number of KB (2 ¹⁰ bytes) of data transferred for write operations, for a particular component over a particular time interval.
Total KB Count	021	SMI-S BSP: volume, controller, storage system	Total number of KB (2 ¹⁰ bytes) of data transferred for read and write operations, for a particular component over a particular time interval.
Read Data Count	022	DS8000: volume, array, controller, storage system	Number of 128 KB chunks (128*2 ¹⁰ bytes) of data transferred for read operations, for a particular component over a particular time interval.
Write Data Count	023	DS8000: volume, array, controller, storage system	Number of 128 KB chunks (128*2 ¹⁰ bytes) of data transferred for write operations, for a particular component over a particular time interval.
Read Sector Count	071	IBM Spectrum Virtualize: node	Number of cache sectors read, for a particular component over a particular time interval.
Write Sector Count	072	IBM Spectrum Virtualize: node	Number of cache sectors written, for a particular component over a particular time interval.
Fast Write Mode Sector Count	074	IBM Spectrum Virtualize: node	Number of written cache sectors processed in fast-write write mode, for a particular component over a particular time interval. Included as part of the "Write Sector Count".
Write Through Mode Sector Count	075	IBM Spectrum Virtualize: node	Number of written cache sectors processed in write-through write mode, for a particular component over a particular time interval. Included as part of the "Write Sector Count".
Flush Through Mode Sector Count	076	IBM Spectrum Virtualize: node	Number of written cache sectors processed in flush-through write mode, for a particular component over a particular time interval. Included as part of the "Write Sector Count".
Cache Overflow Mode Sector Count	077	IBM Spectrum Virtualize: node	Number of written cache sectors processed in cache overflow write mode, for a particular component over a particular time interval. Included as part of the "Write Sector Count".
Non-Preferred Node KB Count	097	IBM Spectrum Virtualize: I/O group, volume	Number of KB (both read and write) that were transferred using the non-preferred nodes of the volumes. This value is a subset of the total KB transferred for the component, and is normally expected to be 0.

Counter	Counter type	Resource or component type	Description
Read Hit KB Percentage for Small Transfers	120	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of read operations that resulted in cache hits. Applies to data transfers ≤ 8 KB.
Read Hit KB Percentage for Medium Transfers	121	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of read operations that resulted in cache hits. Applies to data transfers > 8 KB and ≤ 64 KB.
Read Hit KB Percentage for Large Transfers	122	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of read operations that resulted in cache hits. Applies to data transfers > 64 KB and ≤ 512 KB.
Read Hit KB Percentage for Huge Transfers	123	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of read operations that resulted in cache hits. Applies to data transfers > 512 KB.
Read Miss KB Percentage for Small Transfers	124	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of read operations that resulted in cache misses. Applies to data transfers ≤ 8 KB.
Read Miss KB Percentage for Medium Transfers	125	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of read operations that resulted in cache misses. Applies to data transfers > 8 KB and ≤ 64 KB.
Read Miss KB Percentage for Large Transfers	126	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of read operations that resulted in cache misses. Applies to data transfers > 64 KB and ≤ 512 KB.
Read Miss KB Percentage for Huge Transfers	127	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of read operations that resulted in cache misses. Applies to data transfers > 512 KB.
Write Hit KB Percentage for Small Transfers	128	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of write operations that resulted in cache hits. Applies to data transfers ≤ 8 KB.
Write Hit KB Percentage for Medium Transfers	129	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of write operations that resulted in cache hits. Applies to data transfers > 8 KB and ≤ 64 KB.
Write Hit KB Percentage for Large Transfers	130	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of write operations that resulted in cache hits. Applies to data transfers > 64 KB and ≤ 512 KB.
Write Hit KB Percentage for Huge Transfers	131	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of write operations that resulted in cache hits. Applies to data transfers > 512 KB.
Write Miss KB Percentage for Small Transfers	132	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of write operations that resulted in cache misses. Applies to data transfers ≤ 8 KB.
Write Miss KB Percentage for Medium Transfers	133	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of write operations that resulted in cache misses. Applies to data transfers > 8 KB and ≤ 64 KB.
Write Miss KB Percentage for Large Transfers	134	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of write operations that resulted in cache misses. Applies to data transfers > 64 KB and ≤ 512 KB.
Write Miss KB Percentage for Huge Transfers	135	XIV: volume, module, storage system	Percentage of KB of data that is transferred over a specified interval because of write operations that resulted in cache misses. Applies to data transfers > 512 KB.
Read SSD Data Hit Percentage for Small Transfer Sizes	157	XIV: volume, pool, module, host connection, storage system	Percentage of data transfers over a specified time interval that accessed SSD cache memory. Applies to data transfers ≤ 8 KB.
Read SSD Data Hit Percentage for Medium Transfer Sizes	158	XIV: volume, pool, module, host connection, storage system	Percentage of data transfers over a specified time interval that accessed SSD cache memory. Applies to data transfers > 8 KB and ≤ 64 KB.
Read SSD Data Hit Percentage for Large Transfer Sizes	159	XIV: volume, pool, module, host connection, storage system	Percentage of data transfers over a specified time interval that accessed SSD cache memory. Applies to data transfers > 64 KB and ≤ 512 KB.
Read SSD Data Hit Percentage for Huge Transfer Sizes	160	XIV: volume, pool, module, host connection, storage system	Percentage of data transfers over a specified time interval that accessed SSD cache memory. Applies to data transfers > 512 KB.
Compressed Volumes KB Count	183	IBM Spectrum Virtualize: node, I/O group, storage system	The total number of KiB that were read from or written to compressed volumes.

Table 4. Volume service time counters

Counter	Counter type	Resource or component type	Description
Read Service Time	025	IBM Spectrum Virtualize: volume, node, I/O group, pool, storage system XIV: volume, module, storage system	Number of milliseconds that it took to service all read operations, for a particular component over a particular time interval.
Write Service Time	026	IBM Spectrum Virtualize: volume, node, I/O group, pool, storage system XIV: volume, module, storage system	Number of milliseconds that it took to service all write operations, for a particular component over a particular time interval.
Read Service Periods	028	DS8000: volume, array, controller, storage system SMI-S BSP: volume, controller, storage system	Number of time periods that it took to service all read operations, for a particular component over a particular time interval. The exact time interval length is device-dependent, but is typically 16 milliseconds for DS devices.
Write Service Periods	029	DS8000: volume, array, controller, storage system SMI-S BSP: volume, controller, storage system	Number of time periods that it took to service all write operations, for a particular component over a particular time interval. The exact time interval length is device-dependent, but is typically 16 milliseconds for DS devices.

Counter	Counter type	Resource or component type	Description
Overall Service Periods	030	SMI-S BSP: volume, controller, storage system	Total number of time periods that it took to service all I/O operations (read and write), for a particular component over a particular time interval. The exact time interval length is device-dependent, but is typically 16 milliseconds for DS devices.
Peak Read Response Time	088	IBM Spectrum Virtualize: volume, node, I/O group, storage system	The peak (worst) response time for disk reads for a particular component over a particular time interval. Value in microseconds.
Peak Write Response Time	089	IBM Spectrum Virtualize: volume, node, I/O group, storage system	The peak (worst) response time for disk writes for a particular component over a particular time interval. Value in microseconds.
Global Mirror Service Time	092	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of additional milliseconds it took to service all secondary write operations for Global Mirror, beyond the time required to service the primary writes, for a particular component over a particular time interval.
Overall Host Attributed Service Time	095	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Cumulative transfer latency, in microseconds, coalesced for both reads and writes excluding any status-only transfers. This value is the time taken for hosts to respond to transfer-ready notifications from the nodes (for read) and the time taken for hosts to send the write data after receiving transfer-ready notification from the nodes (for write). This value is provided as an aid to diagnose slow hosts and poorly performing fabrics.
Read Cache Hit Service Periods	099	BSP: volume, controller, storage system	Number of time periods that it took to service all read cache hit operations, for a particular component over a particular time interval. The exact time interval length is device-dependent.
Write Cache Hit Service Periods	100	BSP: volume, controller, storage system	Number of time periods that it took to service all write cache hit operations, for a particular component over a particular time interval. The exact time interval length is device-dependent.
Idle Periods	102	BSP: volume, storage system	The number of time periods during which the component was idle. The exact time interval length is device-dependent.
Maintenance Service Periods	103	BSP: volume, controller, storage system	Number of time periods it took to service all disk maintenance operations, for a particular component over a particular time interval. The exact time interval length is device-dependent.
Read Hit Time Percentage for Small Transfers	136	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on read operations that resulted in cache hits. Applies to data transfers ≤ 8 KB.
Read Hit Time Percentage for Medium Transfers	137	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on read operations that resulted in cache hits. Applies to data transfers > 8 KB and ≤ 64 KB.
Read Hit Time Percentage for Large Transfers	138	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on read operations that resulted in cache hits. Applies to data transfers > 64 KB and ≤ 512 KB.
Read Hit Time Percentage for Huge Transfers	139	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on read operations that resulted in cache hits. Applies to data transfers > 512 KB.
Read Miss Time Percentage for Small Transfers	140	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on read operations that resulted in cache misses. Applies to data transfers ≤ 8 KB.
Read Miss Time Percentage for Medium Transfers	141	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on read operations that resulted in cache misses. Applies to data transfers > 8 KB and ≤ 64 KB.
Read Miss Time Percentage for Large Transfers	142	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on read operations that resulted in cache misses. Applies to data transfers > 64 KB and ≤ 512 KB.
Read Miss Time Percentage for Huge Transfers	143	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on read operations that resulted in cache misses. Applies to data transfers > 512 KB.
Write Hit Time Percentage for Small Transfers	144	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on write operations that resulted in cache hits. Applies to data transfers ≤ 8 KB.
Write Hit Time Percentage for Medium Transfers	145	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on write operations that resulted in cache hits. Applies to data transfers > 8 KB and ≤ 64 KB.
Write Hit Time Percentage for Large Transfers	146	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on write operations that resulted in cache hits. Applies to data transfers > 64 KB and ≤ 512 KB.
Write Hit Time Percentage for Huge Transfers	147	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on write operations that resulted in cache hits. Applies to data transfers > 512 KB.
Write Miss Time Percentage for Small Transfers	148	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on write operations that resulted in cache misses. Applies to data transfers ≤ 8 KB.

Counter	Counter type	Resource or component type	Description
Write Miss Time Percentage for Medium Transfers	149	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on write operations that resulted in cache misses. Applies to data transfers > 8 KB and ≤ 64 KB.
Write Miss Time Percentage for Large Transfers	150	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on write operations that resulted in cache misses. Applies to data transfers > 64 KB and ≤ 512 KB.
Write Miss Time Percentage for Huge Transfers	151	XIV: volume, module, storage system	Percentage of service time, over a specified interval, that was spent on write operations that resulted in cache misses. Applies to data transfers > 512 KB.
Read SSD Response Time for Small Transfer Sizes	161	XIV: volume, pool, module, host connection, storage system	Average number of milliseconds that was required to service each hit operation of SSD read memory over a specified time interval. Applies to transfer operations ≤ 8 KB.
Read SSD Response Time for Medium Transfer Sizes	162	XIV: volume, pool, module, host connection, storage system	Average number of milliseconds that was required to service each hit operation of SSD read memory over a specified time interval. Applies to transfer operations > 8 KB and ≤ 64 KB.
Read SSD Response Time for Large Transfer Sizes	163	XIV: volume, pool, module, host connection, storage system	Average number of milliseconds that was required to service each hit operation of SSD read memory over a specified time interval. Applies to transfer operations > 64 KB and ≤ 512 KB.
Read SSD Response Time for Huge Transfer Sizes	164	XIV: volume, pool, module, host connection, storage system	Average number of milliseconds that was required to service each hit operation of SSD read memory over a specified time interval. Applies to transfer operations > 512 KB.
Compressed Volumes Service Time	184	IBM Spectrum Virtualize: node, I/O group, storage system	The total number of milliseconds that it took to service all read and write operations for compressed volumes.
Cache Prestage Service Time - Volume Cache	618	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of milliseconds that it took to service all prestage I/Os measured at the volume copy cache.
Cache Stage Service Time - Volume Cache	620	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of milliseconds that it took to complete all stage operations in the volume cache. That is, the time that it took to do read operations from the disk to the volume cache.
Cache Stage Service Time - Volume Copy Cache	619	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of milliseconds that it took to complete all stage operations in the volume copy cache. That is, the time that it took to do read operations from the disk to the volume copy cache.
Cache Destage Service Time - Volume Cache	622	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of milliseconds that it took to complete all destage operations in the volume cache. That is, the time that it took to do write operations from the volume cache to the disk.
Cache Destage Service Time - Volume Copy Cache	621	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of milliseconds that it took to complete all destage operations in the volume copy cache. That is, the time that it took to do write operations from the volume copy cache to the disk.
Cache Data Transfer Service Time	179	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of milliseconds that is taken to transfer a track from the cache to the host, including any queuing time that occurs because of throttling. This is similar to the Host Attributed Service Time but includes any queuing delay incurred at the host layer.
Exclusive Track Lock Latency - Volume Cache	638	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The average number of microseconds that it took to acquire an exclusive track lock at the volume cache.
Exclusive Track Lock Latency - Volume Copy Cache	637	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The average number of microseconds that it took to acquire an exclusive track lock at the volume copy cache.

Counter	Counter type	Resource or component type	Description
Shared Track Lock Latency - Volume Cache	640	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The average number of microseconds that it took to acquire a shared track lock at the volume cache.
Shared Track Lock Latency - Volume Copy Cache	639	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The average number of microseconds that it took to acquire a shared track lock at the volume copy cache.
Cache Queue Time IOCB - Volume Cache	642	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume cache spent queuing for I/O control blocks.
Cache Queue Time IOCB - Volume Copy Cache	641	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume copy cache spent queuing for I/O control blocks.
Cache Queue Time CTCB - Volume Cache	644	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume cache spent queuing for cache track control blocks.
Cache Queue Time CTCB - Volume Copy Cache	643	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume copy cache spent queuing for cache track control blocks.
Cache Queue Time PNRC - Volume Cache	646	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume cache spent queuing for remote credits on a preferred node.
Cache Queue Time PNRC - Volume Copy Cache	645	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume copy cache spent queuing for remote credits on a preferred node.
Cache Queue Time NNRC - Volume Cache	648	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume cache spent queuing for remote credits on a non-preferred node.
Cache Queue Time NNRC - Volume Copy Cache	647	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume copy cache spent queuing for remote credits on a non-preferred node.
Cache Queue Time ARC - Volume Cache	650	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume cache spent queuing for remote credits for administrative tasks.
Cache Queue Time ARC - Volume Copy Cache	649	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume copy cache spent queuing for remote credits for administrative tasks.
Cache Queue Time CDCB - Volume Cache	652	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume cache spent queuing for control blocks in the cache directory.
Cache Queue Time CDCB - Volume Copy Cache	651	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume copy cache spent queuing for control blocks in the cache directory.
Cache Queue Time CBR - Volume Cache	654	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume cache spent queuing for buffer resources.
Cache Queue Time CBR - Volume Copy Cache	653	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume copy cache spent queuing for buffer resources.
Cache Queue Time CHR - Volume Cache	656	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume cache spent queuing for hardening rights.
Cache Queue Time CHR - Volume Copy Cache	655	IBM Spectrum Virtualize ¹ : node	The average number of microseconds that the volume copy cache spent queuing for hardening rights.
Note:			
1. This counter is only available for resources that are running IBM Spectrum Virtualize 7.3 or later.			

Table 5. Volume cache counters

Counter	Counter type	Resource or component type	Description
---------	--------------	----------------------------	-------------

Counter	Counter type	Resource or component type	Description
Disk to Cache Transfer Count (normal)	031	DS8000: volume, array, controller, storage system	Number of tracks transferred from disk to cache, for nonsequential I/O requests, for a particular component over a particular time interval.
Disk to Cache Transfer Count (sequential)	032	DS8000: volume, array, controller, storage system	Number of tracks transferred from disk to cache, for sequential I/O requests, for a particular component over a particular time interval.
Cache to Disk Transfer Count	033	DS8000: volume, array, controller, storage system	Number of tracks transferred from cache to disk, for both sequential and nonsequential I/O requests, for a particular component over a particular time interval.
Prestage Track Count	078	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of prestage reads (tracks of read-ahead data) initiated for a particular component over a particular time interval.
Destage Track Count	079	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of destage writes (tracks of dirty write data) initiated for a particular component over a particular time interval.
Prestage Tracks Cache Hits	080	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Number of read track cache hits that were serviced from prestaged tracks for a particular component over a particular time interval. Included as part of the "Read Track Cache Hits".
Destage Tracks Cache Hits	081	IBM Spectrum Virtualize: volume, node, I/O group, storage system	The number of write track cache hits for a particular component over a particular time interval. That is, the track writes that modified only data that was already marked "dirty" in the cache.
Prestage Sector Count	082	IBM Spectrum Virtualize: node	Number of prestage sectors (sectors of read-ahead data) initiated for a particular component over a particular time interval.
Destage Sector Count	083	IBM Spectrum Virtualize: node	Number of destage writes (sectors of dirty write data) initiated for a particular component over a particular time interval.
Prestage Tracks Cache Hits Sector Count	084	IBM Spectrum Virtualize: node	Number of sectors read for cache hits that were serviced from prestaged tracks for a particular component over a particular time interval. Included as part of the "Read Track Cache Hits Sector Count".
Destage Tracks Cache Hits Sector Count	085	IBM Spectrum Virtualize: node	The number of sectors written for cache hits for a particular component over a particular time interval. That is, the writes that modified only data that was already marked "dirty" in the cache.
Read Sector Count - Volume Cache	62	IBM Spectrum Virtualize: volume, pool, node, I/O group, host connection, storage system	The number of cache sectors or tracks read by the volume cache. For resources that are running a version of IBM Spectrum Virtualize that is earlier than 7.3, the unit of the counter is tracks.
Read Sector Count - Volume Copy Cache	601	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache sectors read by the volume copy cache.
Write Sector Count - Volume Cache	63	IBM Spectrum Virtualize: volume, pool, node, I/O group, host connection, storage system	The number of cache sectors or tracks written by the volume cache. For resources that are running a version of IBM Spectrum Virtualize that is earlier than 7.3, the unit of the counter is tracks.
Write Sector Count - Volume Copy Cache	602	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache sectors written by the volume copy cache.
Fast Write Mode Count - Volume Cache	65	IBM Spectrum Virtualize: volume, pool, node, I/O group, host connection, storage system	The number of cache sectors or tracks written in fast-write mode in the volume cache. For resources that are running a version of IBM Spectrum Virtualize that is earlier than 7.3, the unit of the counter is tracks.
Fast Write Mode Count - Volume Copy Cache	603	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache sectors written in fast-write mode in the volume copy cache.
Flush-through Mode Count - Volume Cache	67	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache sectors or tracks written in flush-through write mode in the volume cache. For resources that are running a version of IBM Spectrum Virtualize that is earlier than 7.3, the unit of the counter is tracks.
Flush-through Mode Count - Volume Copy Cache	604	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache sectors written in flush-through write mode in the volume copy cache.
Write-through Mode Count - Volume Cache	66	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache sectors or tracks written in write-through write mode in the volume cache. For resources that are running a version of IBM Spectrum Virtualize that is earlier than 7.3, the unit of the counter is tracks.
Write-through Mode Count - Volume Copy Cache	605	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache sectors written in write-through write mode in the volume copy cache.
Cache Overflow Mode Track Count	068	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache track that were processed in Cache-Overflow write mode.
Prestage Count	78	IBM Spectrum Virtualize: volume, pool, node, I/O group, host connection, storage system	The number of cache sectors or tracks prestaged into the volume copy cache. For resources that are running a version of IBM Spectrum Virtualize that is earlier than 7.3, the unit of the counter is tracks.
Destage Count - Volume Cache	79	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache sectors or tracks destaged from the volume cache. For resources that are running a version of IBM Spectrum Virtualize that is earlier than 7.3, the unit of the counter is tracks.
Destage Count - Volume Copy Cache	606	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of cache sectors destaged from the volume copy cache.
Cache Data Transfer Count	178	IBM Spectrum Virtualize ¹ : volume, pool, node, I/O group, host connection, storage system	The number of tracks transferred from the volume cache to the host I/O layer. In combination with the Cache Data Transfer Service Time counter determines the Data Transfer Response Time.

Counter	Counter type	Resource or component type	Description
Note:			
1. This counter is only available for resource that are running IBM Spectrum Virtualize 7.3 or later.			

Table 6. Volume miscellaneous counters

Counter	Counter type	Resource or component type	Description
NVS Allocation Count	034	DS8000: volume	Number of NVS space allocations for a particular component over a particular time interval.
DFW I/O Count (normal)	035	DS8000: volume, array, controller, storage system	Number of nonsequential "DASD Fast Write" requests for a particular component over a particular time interval. This value typically applies only for count-key data (CKD).
DFW I/O Count (sequential)	036	DS8000: volume, array, controller, storage system	Number of sequential "DASD Fast Write" requests for a particular component over a particular time interval. This value typically applies only for CKD.
DFW Delayed I/O Count	037	DS8000: volume, array, controller, storage system	Number of I/O requests that were delayed due to nonvolatile storage (NVS) space constraints, for a particular component over a particular time interval.
Cache Delayed I/O Count	038	DS8000: volume	Number of I/O requests that were delayed due to cache space constraints, for a particular component over a particular time interval.
Record Mode Read I/O Count	039	DS8000: volume, array, controller, storage system	Number of record-mode read I/O operations for a particular component over a particular time interval.
Record Mode Read Cache Hits	040	DS8000: volume, array, controller, storage system	Number of cache hits for record-mode read I/O operations for a particular component over a particular time interval.
Quick Write Promote Count	041	DS8000: volume, array, controller, storage system	Number of quick write promote operations for a particular component over a particular time interval. This value typically applies only for CKD.
CFW Read I/O Count	042	DS8000: volume	Number of "Cache Fast Write" read I/O requests for a particular component over a particular time interval. This value typically applies only for CKD.
CFW Write I/O Count	043	DS8000: volume	Number of "Cache Fast Write" write I/O requests for a particular component over a particular time interval. This value typically applies only for CKD.
CFW Read Hits	044	DS8000: volume	Number of cache hits for "Cache Fast Write" read I/O requests for a particular component over a particular time interval. This value typically applies only for CKD.
CFW Write Hits	045	DS8000: volume	Number of cache hits for "Cache Fast Write" write I/O requests for a particular component over a particular time interval. This value typically applies only for CKD.
XRC Read Count	046	DS8000: volume	Number of track reads from the Concurrent Copy or XRC Sidefile for a particular component over a particular time interval. This value typically applies only for CKD.
XRC Contaminating Write Count	047	DS8000: volume	Number of contaminating writes for Concurrent Copy or XRC operations for a particular component over a particular time interval. This value typically applies only for CKD.
PPRC Transfer Count	048	DS8000: volume, array, controller, storage system	Number of tracks or portions of tracks that were transferred to the secondary device of a Peer-to-Peer Remote Copy (PPRC) pair, for a particular component over a particular time interval. This counter shows the activity for the source of the PPRC relationship, but shows no activity for the target.
ITA I/O Count	049	DS8000: volume	Number of irregular track access I/O requests, for a particular component over a particular time interval. This value typically applies only for CKD.
ITA Hits	050	DS8000: volume	Number of cache hits for irregular track access I/O requests, for a particular component over a particular time interval. This value typically applies only for CKD.
Rank Read I/O Count	051	DS8000: volume	Number of read I/O operations issued to the physical arrays associated with a particular component over a particular time interval.
Rank Write I/O Count	052	DS8000: volume	Number of write I/O operations issued to the physical arrays associated with a particular component over a particular time interval.
Rank Read Data Chunk Count	053	DS8000: volume	Number of 128 KB chunks (128*2 ¹⁰ bytes) of data transferred for read I/O operations issued to the physical arrays associated with a particular component over a particular time interval.
Rank Write Data Chunk Count	054	DS8000: volume	Number of 128 KB chunks (128*2 ¹⁰ bytes) of data transferred for write I/O operations issued to the physical arrays associated with a particular component over a particular time interval.
Rank Read Service Periods	055	DS8000: volume	Number of 16 millisecond time periods that it took to service all read operations issued to the physical arrays associated with a particular component over a particular time interval.
Rank Write Service Periods	056	DS8000: volume	Number of 16 millisecond time periods that it took to service all write operations issued to the physical arrays associated with a particular component over a particular time interval.
ICL Read I/O Count	057	DS8000: volume	Number of Inhibit Cache Loading I/O requests, for a particular component over a particular time interval. This value typically applies only for CKD.
BC Write I/O Count	058	DS8000: volume	Number of Bypass Cache I/O requests, for a particular component over a particular time interval. This value typically applies only for CKD.
Rank Transfer Time	059	DS8000: volume	Number of milliseconds of lower-interface I/O activity, for a particular component over a particular time interval. This value typically applies only for CKD.
Read Cache KB Count	086	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Amount of read cache, in KB, that was in use for a particular component, at the start of the particular time interval.

Counter	Counter type	Resource or component type	Description
Write Cache KB Count	087	IBM Spectrum Virtualize: volume, node, I/O group, storage system	Amount of write cache, in KB, that was in use for a particular component, at the start of the particular time interval.

Table 7. Volume special computed values

Counter	Counter type	Resource or component type	Description
Cache Holding Time	060	DS8000: controller, storage system	Average cache holding time, in seconds, for I/O data in this subsystem controller (cluster). Shorter time periods indicate adverse performance.
Utilization Percentage	061	DS8000: volume, array, controller, storage system IBM Spectrum Virtualize: volume, node, I/O group, storage system XIV: volume	Approximate utilization percentage for a particular component over a particular time interval. The value is multiplied by 100 to gain 2 extra digits of precision; for example, 3860 would indicate that the component is running at approximately 38.6% of capacity. The value might exceed 100% to indicate stress beyond the capacity of the component.
Compression CPU Utilization	175	IBM Spectrum Virtualize: node, I/O group, storage system	The approximate percentage of time that the processors used for data compression I/O tasks are busy.
System CPU Core Utilization	176 (Core 1 to 8) 185 (Core 9 to 16) 186 (Core 17 to 24) 187 (Core 25 to 32)	IBM Spectrum Virtualize: node, I/O group, storage system	The approximate percentage of time that each set of eight processor cores was busy with system I/O tasks. The value is encoded such that each byte contains the percentage of one of the cores; the low-order byte is for core 1.
Compression CPU Core Utilization	177 (Core 1 to 8) 188 (Core 9 to 16) 189 (Core 17 to 24) 190 (Core 25 to 32)	IBM Spectrum Virtualize: node, I/O group, storage system	The approximate percentage of time that each set of eight processor cores was busy with data compression tasks. The value is encoded such that each byte contains the percentage of one of the cores; the low-order byte is for core 1.

Disk counters

Performance counters for back-end storage systems are divided into the following categories:

- [Back-end storage system I/O counters](#)
- [Back-end storage system data counters](#)
- [Back-end storage system service time counters](#)
- [Other back-end counters for storage systems](#)

Table 8. Back-end storage system I/O counters

Counter	Counter type	Resource or component type	Description
Backend Read I/O Count	201	DS8000: rank, array, controller, storage system IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system	Number of I/O operations for read requests, for a particular component over a particular time interval.
Backend Write I/O Count	202	DS8000: rank, array, controller, storage system IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system	Number of I/O operations for write requests, for a particular component over a particular time interval.

Table 9. Back-end storage systems data counters

Counter	Counter type	Resource or component type	Description
Backend Read KB Count	204	IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system IBM FlashSystem® 900: flash module, storage system	Number of KB (2 ¹⁰ bytes) of data that were transferred for read operations, for a particular component over a particular time interval.
Backend Write KB Count	205	IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system IBM FlashSystem 900: flash module, storage system	Number of KB (2 ¹⁰ bytes) of data that were transferred for write operations, for a particular component over a particular time interval.
Backend Read Data Count	207	DS8000: rank, array, controller, storage system	Number of 128 KB chunks (128*2 ¹⁰ bytes) of data that were transferred for read operations, for a particular component over a particular time interval.
Backend Write Data Count	208	DS8000: rank, array, controller, storage system	Number of 128 KB chunks (128*2 ¹⁰ bytes) of data that were transferred for write operations, for a particular component over a particular time interval.

Table 10. Back-end service time counters

Counter	Counter type	Resource or component type	Description
---------	--------------	----------------------------	-------------

Counter	Counter type	Resource or component type	Description
Backend Read Service Time	210	IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system	Number of milliseconds that it took to service all read operations, for a particular component over a particular time interval. For IBM Spectrum Virtualize, this value is the <i>external</i> response time of the managed disk.
Backend Write Service Time	211	IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system	Number of milliseconds that it took to service all write operations, for a particular component over a particular time interval. For IBM Spectrum Virtualize, this value is the <i>external</i> response time of the managed disk.
Backend Read Service Periods	213	DS8000: rank, array, controller, storage system	Number of 16 millisecond time periods that it took to service all read operations, for a particular component over a particular time interval.
Backend Write Service Periods	214	DS8000: rank, array, controller, storage system	Number of 16 millisecond time periods that it took to service all write operations, for a particular component over a particular time interval.
Backend Read Queued Service Time	216	IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system	Number of milliseconds that it took to service all read operations, for a particular component over a particular time interval. Includes the time that pending I/O operations spent on the queue waiting to be issued to the backend disk device.
Backend Write Queued Service Time	217	IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system	Number of milliseconds that it took to service all write operations, for a particular component over a particular time interval. Includes the time that pending I/O operations spent on the queue waiting to be issued to the backend disk device.
Peak Backend Read Response Time	219	IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system	The peak (worst) response time for disk reads for a particular component over a particular time interval, in microseconds. For IBM Spectrum Virtualize, this value is the peak <i>external</i> response time of the managed disks.
Peak Backend Write Response Time	220	IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system	The peak (worst) response time for disk writes for a particular component over a particular time interval, in microseconds. For IBM Spectrum Virtualize, this is the peak <i>external</i> response time of the managed disks.
Peak Backend Read Queued Response Time	222	IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system	The peak (worst) response time for disk reads for a particular component over a particular time interval, in microseconds. Includes the time that pending I/O operations spent on the queue waiting to be issued to the backend disk device.
Peak Backend Write Queued Response Time	223	IBM Spectrum Virtualize: managed disk, pool, node, I/O group, storage system	The peak (worst) response time for disk reads for a particular component over a particular time interval, in microseconds. Includes the time that pending I/O operations spent on the queue waiting to be issued to the backend disk device.

Table 11. Other back-end counters for storage systems

Counter	Counter type	Resource or component type	Description
Flash Health Percentage	238	Flash module, storage system	The overall flash module health. The health percentage is calculated based on the number of unusable blocks on the flash module.

Storage system and switch port counters

Performance counters for ports and fabrics are divided into the following categories:

- [Storage system and switch port I/O or frame counters](#)
- [Storage system and switch port data counters](#)
- [Storage system port service time counters](#)
- [Storage system and switch port utilization percentages](#)
- [Storage system and switch port error counts](#)

Table 12. Storage system and switch port I/O or frame counters

Counter	Counter type	Resource or component type	Description
Port Send I/O Count	401	DS8000: port, storage system IBM FlashSystem 900: port XIV: port	Number of read I/O operations processed (or write I/O operations initiated) by a particular component over a particular time interval. For ports, available only for hourly/daily data.
Port Receive I/O Count	402	DS8000: port, storage system IBM FlashSystem 900: port XIV: port	Number of write I/O operations processed (or read operations initiated) by a particular component over a particular time interval. For ports, available only for hourly/daily data.
Total Port I/O Count	403	SMI-S BSP: port	Total number of read and write I/O operations processed by a particular component over a particular time interval.
Port Send Frame Count	404	Switch port Outer-switch connection ¹ Switch	Number of frames sent by a particular component over a particular time interval.
Port Receive Frame Count	405	Switch port Inter-switch connection ¹ Switch	Number of frames received by a particular component over a particular time interval.

Counter	Counter type	Resource or component type	Description
Total Port Frame Count	406	Switch port Inter-switch connection ¹ Switch	Number of frames sent and received by a particular component over a particular time interval.
Port FB Send I/O Count	407	DS8000: port	Number of read I/O operations processed (or write operations initiated) for FB requests, for a particular port over a particular time interval.
Port FB Receive I/O Count	408	DS8000: port	Number of write I/O operations processed (or read operations initiated) for FB requests, for a particular port over a particular time interval.
Port CKD Send I/O Count	410	DS8000: port	Number of read I/O operations processed (or write operations initiated) for CKD requests, for a particular port over a particular time interval.
Port CKD Receive I/O Count	411	DS8000: port	Number of write I/O operations processed (or read operations initiated) for CKD requests, for a particular port over a particular time interval.
Port PPRC Send I/O Count	413	DS8000: port, storage system	Number of write operations initiated by the PPRC primary, for a particular port over a particular time interval.
Port PPRC Receive I/O Count	414	DS8000: port, storage system	Number of write operations processed by the PPRC secondary, for a particular port over a particular time interval.
Port to Host Send I/O Count	442	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of exchanges (I/Os) initiated to host computers from a particular port over a particular time interval.
Port to Host Receive I/O Count	443	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of exchanges (I/Os) received from host computers by a particular port over a particular time interval.
Port to Disk Send I/O Count	445	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of exchanges (I/Os) initiated to storage systems from a particular port over a particular time interval.
Port to Disk Receive I/O Count	446	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of exchanges (I/Os) received from storage systems by a particular port over a particular time interval.
Port to Local Send I/O Count	448	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of exchanges (I/Os) initiated to other IBM Spectrum Virtualize nodes on the local IBM Spectrum Virtualize cluster from a particular port over a particular time interval.
Port to Local Receive I/O Count	449	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of exchanges (I/Os) received from other IBM Spectrum Virtualize nodes on the local cluster by a particular port over a particular time interval.
Port to Remote Send I/O Count	451	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of exchanges (I/Os) initiated to a remote IBM Spectrum Virtualize cluster from a particular port over a particular time interval.
Port to Remote Receive I/O Count	452	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of exchanges (I/Os) received from a remote cluster by a particular port over a particular time interval.
Port to Local Send Message Count	478	IBM Spectrum Virtualize: node, I/O group, storage system	Number of messages sent to nodes on the local IBM Spectrum Virtualize cluster, by the ports associated with a particular component, over a particular time interval. This count might be larger or smaller than the Port to Local Send I/O Count, because multiple small messages might be batched into a single I/O, or a single large message split between multiple I/Os.
Port to Local Receive Message Count	479	IBM Spectrum Virtualize: node, I/O group, storage system	Number of messages received from nodes on the local IBM Spectrum Virtualize cluster, by the ports associated with a particular component, over a particular time interval. This count might be larger or smaller than the Port to Local Receive I/O Count, because multiple small messages might be batched into a single I/O, or a single large message split between multiple I/Os.
Port to Remote Send Message Count	481	IBM Spectrum Virtualize: node, I/O group, storage system	Number of messages sent to nodes on the remote IBM Spectrum Virtualize cluster, by the ports associated with a particular component, over a particular time interval. This count might be larger or smaller than the Port to Remote Send I/O Count, because multiple small messages might be batched into a single I/O, or a single large message split between multiple I/Os.
Port to Remote Receive Message Count	482	IBM Spectrum Virtualize: node, I/O group, storage system	Number of messages received from nodes on the remote IBM Spectrum Virtualize cluster, by the ports associated with a particular component, over a particular time interval. This count might be larger or smaller than the Port to Remote Receive I/O Count, because multiple small messages might be batched into a single I/O, or a single large message split between multiple I/Os.
<p>Note:</p> <p>1. Performance data for inter-switch connections contains performance data for the following resources:</p> <ul style="list-style-type: none"> Switch ports on ISLs and NPV links. Switch trunks for ISL trunks, ICL trunks, F_port channels, and port channels. <p>The performance data for both resources is displayed together in the IBM Spectrum® Control GUI. However, to obtain this data from the CLI, you must run separate queries.</p>			

Table 13. Storage system and switch port data counters

Counter	Counter type	Resource or component type	Description
Port Send KB Count	416	DS8000: storage system IBM FlashSystem 900: port Switch port Inter-switch connection ¹ Switch XIV: port	Number of KB (2^{10} bytes) of data that were sent from a particular component over a particular time interval. This value is read KB if the port is the target of a read operation, or write KB if the port is the initiator of a write operation (for example, a PPRC primary).
Port Receive KB Count	417	DS8000: storage system IBM FlashSystem 900: port Switch port Inter-switch connection ¹ Switch XIV: port	Number of KB (2^{10} bytes) of data that were received by a particular component over a particular time interval. This value is write KB if the port is the target of a write operation, or read KB if the port is the initiator of a read operation.
Total Port KB Count	418	Switch port Inter-switch connection ¹ Switch SMI-S BSP: port	Total number of KB (2^{10} bytes) of data that were transferred, by a particular port over a particular time interval.
Port Peak Send Data Rate ²	419	Switch port Inter-switch connection ¹ Switch	Peak number of megabytes (2^{20} bytes) per second that were sent by a particular port over a time interval.
Port Peak Receive Data Rate	420	Switch port Inter-switch connection ¹ Switch	Peak number of megabytes (2^{20} bytes) per second that were received by a particular port over a time interval.
Port FB Send Data Count	421	DS8000: port	Number of 128 KB chunks (128×2^{10} bytes) of data that were sent for FB requests, from a particular component over a particular time interval. This value is read KB if the port is the target of a read operation, or write KB if the port is the initiator of a write operation (for example, a PPRC primary).
Port FB Receive Data Count	422	DS8000: port	Number of 128 KB chunks (128×2^{10} bytes) of data that were received for FB requests, by a particular component over a particular time interval. This value is write KB if the port is the target of a write operation, or read KB if the port is the initiator of a read operation.
Port CKD Send Data Count	424	DS8000: port	Number of 128 KB chunks (128×2^{10} bytes) of data that were sent for CKD requests, from a particular component over a particular time interval. This value is read KB if the port is the target of a read operation, or write KB if the port is the initiator of a write operation (for example, a PPRC primary).
Port CKD Receive Data Count	425	DS8000: port	Number of 128 KB chunks (128×2^{10} bytes) of data that were received for CKD requests, by a particular component over a particular time interval. This value is write KB if the port is the target of a write operation, or read KB if the port is the initiator of a read operation.
Port PPRC Send Data Count	427	DS8000: port, storage system	Number of 128 KB chunks (128×2^{10} bytes) of data that were sent by a particular port over a particular time interval, as result of write operations initiated by a PPRC primary.
Port PPRC Receive Data Count	428	DS8000: port, storage system	Number of 128 KB chunks (128×2^{10} bytes) of data that were received by a particular port over a particular time interval, as result of write operations processed by a PPRC secondary.
Port to Host Send KB Count	454	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of KB of data sent to host computers from a particular port over a particular time interval.
Port to Host Receive KB Count	455	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of KB of data received from host computers by a particular port over a particular time interval.
Port to Disk Send KB Count	457	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of KB of data sent to storage systems from a particular port over a particular time interval.
Port to Disk Receive KB Count	458	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of KB of data received from storage systems by a particular port over a particular time interval.

Counter	Counter type	Resource or component type	Description
Port to Local Send KB Count	460	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of KB of data sent to other IBM Spectrum Virtualize nodes on the local IBM Spectrum Virtualize cluster from a particular port over a particular time interval.
Port to Local Receive KB Count	461	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of KB of data received from other IBM Spectrum Virtualize nodes on the local cluster by a particular port over a particular time interval.
Port to Remote Send KB Count	463	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of KB of data sent to a remote IBM Spectrum Virtualize cluster from a particular port over a particular time interval.
Port to Remote Receive KB Count	464	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of KB of data received from a remote IBM Spectrum Virtualize cluster by a particular port over a particular time interval.
<p>Note:</p> <ol style="list-style-type: none"> Performance data for inter-switch connections contains performance data for the following resources: <ul style="list-style-type: none"> Switch ports on ISLs and NPV links. Switch trunks for ISL trunks, ICL trunks, F_port channels, and port channels. <p>The performance data for both resources is displayed together in the IBM Spectrum Control GUI. However, to obtain this data from the CLI, you must run separate queries.</p> This counter is available only for ports on Brocade switches. 			

Table 14. Storage system port service time counters

Counter	Counter type	Resource or component type	Description
Port Send Service Time	430	DS8000: storage system IBM FlashSystem 900: port XIV: port	Number of milliseconds it took to service all send operations, for a particular component over a particular time interval. This value is read times, if the port is the target of a read operation, or write times if the port is the initiator of a write operation (for example, a PPRC primary).
Port Receive Service Time	431	DS8000: storage system IBM FlashSystem 900: port XIV: port	Number of milliseconds it took to service all receive operations, for a particular port over a particular time interval. This value is write times, if the port is the target of a write operation, or read times if the port is the initiator of a read operation.
Port FB Send Service Periods	433	DS8000: port	Number of 16 millisecond periods it took to service all send operations for FB requests, for a particular component over a particular time interval. This value is read times, if the port is the target of a read operation, or write times if the port is the initiator of a write operation (for example, a PPRC primary).
Port FB Receive Service Periods	434	DS8000: port	Number of 16 millisecond periods it took to service all receive operations for FB requests, for a particular component over a particular time interval. This value is write times, if the port is the target of a write operation, or read times if the port is the initiator of a read operation.
Port CKD Send Service Periods	436	DS8000: port	Number of 16 millisecond periods it took to service all send operations for CKD requests, for a particular component over a particular time interval. This value is read times, if the port is the target of a read operation, or write times if the port is the initiator of a write operation (for example, a PPRC primary).
Port CKD Receive Service Periods	437	DS8000: port	Number of 16 millisecond periods it took to service all receive operations for CKD requests, for a particular component over a particular time interval. This value is write times, if the port is the target of a write operation, or read times if the port is the initiator of a read operation.
Port PPRC Send Service Periods	439	DS8000: port, storage system	Number of 16 millisecond periods it took to service write requests by a PPRC primary, for a particular port over a particular time interval.
Port PPRC Receive Service Periods	440	DS8000: port, storage system	Number of 16 millisecond periods it took to service write requests by a PPRC secondary, for a particular port over a particular time interval.
Port to Local Send Service Time	466	IBM Spectrum Virtualize: node, I/O group, storage system	Number of milliseconds it took to service all send operations to other nodes on the local IBM Spectrum Virtualize cluster, for a particular component over a particular time interval. For IBM Spectrum Virtualize, this value is the <i>external</i> service time of the transfers.
Port to Local Receive Service Time	467	IBM Spectrum Virtualize: node, I/O group, storage system	Number of milliseconds it took to service all receive operations from other nodes on the local IBM Spectrum Virtualize cluster, for a particular component over a particular time interval. For IBM Spectrum Virtualize, this value is the <i>external</i> service time of the transfers.
Port to Local Send Queued Service Time	469	IBM Spectrum Virtualize: node, I/O group, storage system	Number of milliseconds it took to service all send operations to other nodes on the local IBM Spectrum Virtualize cluster, including the time that pending operations spent on the queue, for a particular component over a particular time interval. For IBM Spectrum Virtualize, this value is the <i>queued</i> service time of the transfers.
Port to Local Receive Queued Service Time	470	IBM Spectrum Virtualize: node, I/O group, storage system	Number of milliseconds it took to service all receive operations from other nodes on the local IBM Spectrum Virtualize cluster, including the time that pending operations spent on the queue, for a particular component over a particular time interval. For IBM Spectrum Virtualize, this value is the <i>queued</i> service time of the transfers.

Counter	Counter type	Resource or component type	Description
Port to Remote Send Service Time	472	IBM Spectrum Virtualize: node, I/O group, storage system	Number of milliseconds it took to service all send operations to nodes on the remote IBM Spectrum Virtualize cluster, for a particular component over a particular time interval. For IBM Spectrum Virtualize, this value is the <i>external</i> service time of the transfers.
Port to Remote Receive Service Time	473	IBM Spectrum Virtualize: node, I/O group, storage system	Number of milliseconds it took to service all receive operations from nodes on the remote IBM Spectrum Virtualize cluster, for a particular component over a particular time interval. For IBM Spectrum Virtualize, this value is the <i>queued</i> service time of the transfers.
Port to Remote Send Queued Service Time	475	IBM Spectrum Virtualize: node, I/O group, storage system	Number of milliseconds it took to service all send operations to nodes on the remote IBM Spectrum Virtualize cluster, including the time that pending operations spent on the queue, for a particular component over a particular time interval. For IBM Spectrum Virtualize, this value is the <i>queued</i> service time of the transfers.
Port to Remote Receive Queued Service Time	476	IBM Spectrum Virtualize: node, I/O group, storage system	Number of milliseconds it took to service all receive operations from nodes on the remote IBM Spectrum Virtualize cluster, including the time that pending operations spent on the queue, for a particular component over a particular time interval. For IBM Spectrum Virtualize, this value is the <i>queued</i> service time of the transfers.

Table 15. Storage system and switch port utilization percentages

Counter	Counter type	Resource or component type	Description
Port Send Utilization Percentage	484	DS8000: port	Approximate percentage of time that the associated port was busy with sending data. The value is multiplied by 100 to give 2 extra digits of precision, for example, 3860 would indicate that the port is busy 38.6% of the time. The port is busy sending data when servicing read I/Os from hosts, and when initiating write I/Os from a PPRC primary.
Port Receive Utilization Percentage	485	DS8000: port	Approximate percentage of time that the associated port was busy with receiving data. The value is multiplied by 100 to give 2 extra digits of precision, for example, 3860 would indicate that the port is busy 38.6% of the time. The port is busy receiving data when servicing non-PPRC write I/Os from hosts, and when servicing write I/Os at a PPRC secondary.
Port Send Bandwidth Percentage	487	DS8000: port IBM FlashSystem 900: port IBM Spectrum Virtualize: port Switch port Inter-switch connection ¹	Approximate percentage of bandwidth utilization for send operations by this port, based on the current negotiated speed of the port. The value is multiplied by 100 to give 2 extra digits of precision, for example, 3860 would indicate that 38.6% of the port bandwidth was being used. The port is busy sending data when servicing read I/Os from hosts, and when initiating write I/Os from a PPRC primary.
Port Receive Bandwidth Percentage	488	DS8000: port IBM FlashSystem 900: port IBM Spectrum Virtualize: port Switch port Inter-switch connection ¹	Approximate percentage of bandwidth utilization for receive operations by this port, based on the current negotiated speed of the port. The value is multiplied by 100 to give 2 extra digits of precision, for example, 3860 would indicate that 38.6% of the port bandwidth was being used. The port is busy receiving data when servicing non-PPRC write I/Os from hosts, and when servicing write I/Os at a PPRC secondary.
Port Link Quality Percentage	490	Switch port Inter-switch connection ¹	The estimated link quality of the switch port. The percentage is based on whether the port is an expansion port (E_port) or a fabric port (F_port), and on the error statistics for the port.

Note:

1. Performance data for inter-switch connections contains performance data for the following resources:

- Switch ports on ISLs and NPV links.
- Switch trunks for ISL trunks, ICL trunks, F_port channels, and port channels.

The performance data for both resources is displayed together in the IBM Spectrum Control GUI. However, to obtain this data from the CLI, you must run separate queries.

Table 16. Storage system and switch port error counts

Counter	Counter type	Resource or component type	Description
Operational Status	501	Switch port Inter-switch connection ¹	The operational status of a port over a time interval.
LIP Count	502	Switch port Inter-switch connection ¹ Switch	Number of LIP events that occurred on the associated arbitrated loops, for a particular component over a particular time interval.

Counter	Counter type	Resource or component type	Description
NOS Count	503	Switch port Inter-switch connection ¹ Switch	Number of NOS events that occurred on the associated switched fabrics, for a particular component over a particular time interval.
Error Frame Count ²	504	DS8000: port, storage system	Number of frames received in error by a particular component over a particular time interval.
Discarded Frame Count	505	Switch port Inter-switch connection ¹ Switch	Number of frames that were lost due to a lack of available host buffers, for a particular component over a particular time interval.
Link Failure Count	506	DS8000: port, storage system IBM Spectrum Virtualize: port, node, I/O group, storage system	The number of link errors that were experienced by a particular component over a particular time interval.
Loss of Sync Count	507	DS8000: port, storage system IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of times that synchronization was lost after the last reset of the device, for a particular component over a particular time interval. Synchronization is assumed lost after a timeout interval identified by the Receiver-Transmitter-Timeout property. This count is part of the Link Error Status Block (LESB).
Loss of Signal Count	508	DS8000: port, storage system IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of times that the signal was lost after the last reset of the device, for a particular component over a particular time interval. This count is part of the Link Error Status Block (LESB).
CRC Error Count	509	DS8000: port, storage system IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of frames received in which the CRC in a frame does not match the CRC computed by the receiver. This count is part of the Link Error Status Block (LESB).
Primitive Sequence Protocol Error Count	510	DS8000: port, storage system IBM Spectrum Virtualize: port, node, I/O group, storage system	Count of primitive sequence protocol errors detected at this port. This count is part of the Link Error Status Block (LESB).
Invalid Transmission Word Count	511	DS8000: port, storage system IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of transmission words that had an 8b10b code violation in one or more of its characters; had a K28.5 in its second, third, or fourth character positions; or was an ordered set that had an incorrect Beginning Running Disparity. This count is part of the Link Error Status Block (LESB).
Short Frame Count ³	512	Switch port Inter-switch connection ¹ Switch	Number of frames received that were shorter than 28 octets. The value of 28 is calculated based on an assumption of 24 header bytes plus 4 CRC bytes. The count does not include SOF/EOF bytes which are not data.
Long Frame Count	513	Switch port Inter-switch connection ¹ Switch	Number of frames received that were longer than 2140 octets. The value of 2140 is calculated based on an assumption of 24 header bytes plus 4 CRC bytes plus 2112 bytes of payload.
Address Error Count	514	Switch port Inter-switch connection ¹ Switch	Count of frames received with unknown addressing. An example is an unknown SID.
Buffer Credit Not Provided Count	515	Switch port Inter-switch connection ¹ Switch	Count of occurrences when all input buffers of a port were full and outbound buffer-to-buffer credit transitioned to zero. There is no credit to provide to other side.
Buffer Credit Not Received Count	516	Switch port Inter-switch connection ¹ Switch	Count of transitions in/out of Bbcredit zero state. The other side is not providing any credit.

Counter	Counter type	Resource or component type	Description
Delimiter Error Count	517	Switch port Inter-switch connection ¹ Switch	Count of invalid frame delimiters received at this port. An example is a frame with a class 2 start and a class 3 at the end.
Encoding Disparity Error Count	518	Switch port Inter-switch connection ¹ Switch	Count of disparity errors received at this port.
Link Reset Transmitted Count	519	DS8000: port, storage system	Count of link resets. This value is the number of LRs transmitted.
Link Reset Received Count	520	DS8000: port, storage system	Count of link resets. This value is the number LRs received.
Discarded Class 3 Frame Count ²	521	Switch port Inter-switch connection ¹ Switch	Class3FramesDiscarded is the count of class 3 frames that were discarded upon reception.
Invalid Ordered Set Count	522	Switch port Inter-switch connection ¹ Switch	The total number of invalid ordered sets received.
F-BSY Frame Count ³	523	Switch port Inter-switch connection ¹ Switch	Number of F_BSY frames generated.
P-BSY Frame Count	524	Switch port Inter-switch connection ¹ Switch	Number of P_BSY frames generated.
F-RJT Frame Count ³	525	Switch port Inter-switch connection ¹ Switch	Number of F_BSY frames generated.
P-RJT Frame Count	526	Switch port Inter-switch connection ¹ Switch	Number of P_BSY frames generated.
Buffer Credit Zero Error Count	527	Switch port Inter-switch connection ¹ Switch	Number of times that frame transmission was blocked by a transmit credit of zero.
Zero Buffer Credit Timer	528	IBM Spectrum Virtualize: port, node, I/O group, storage system	Number of microseconds for which the port has been unable to send frames due to lack of buffer credit since the last node reset.
Out of Order Data Count	529	DS8000: port, storage system	Number of times that an out of order frame is detected.
Out of Order ACK Count	530	DS8000: port, storage system	Number of times that an out of order ACK frame is detected.
Duplicate Frame Count	531	DS8000: port, storage system	Number of times a frame was received that has been detected as previously processed.
Invalid Relative Offset Count	532	DS8000: port, storage system	Number of times that a frame was received with bad relative offset in the frame header.
Sequence Timeout Count	533	DS8000: port, storage system	Number of times the port has detected a timeout condition on receiving sequence initiative for a fibre channel exchange.
Class 3 Send Timeout Frame Count ²	534	Brocade switch port Brocade Inter-switch connection ¹ Brocade switch	Average count of class 3 frames per second that were discarded before transmission because of a timeout condition. The timeout condition occurs while the switch or port waits for buffer credit from the receiving port at the other end of the fibre. When you troubleshoot a SAN, use this count to view port conditions that might slow the performance of the resources to which those ports are connected.

Counter	Counter type	Resource or component type	Description
Class 3 Receive Timeout Frame Count ²	535	Brocade switch port Brocade Inter-switch connection ¹ Brocade switch	Average count of class 3 frames per second that were discarded after reception because of a timeout condition. The timeout condition occurs while the transmitting port waits for buffer credit from a port at the other end of the fibre. When you troubleshoot a SAN, use this count to view port conditions that might slow the performance of the resources to which those ports are connected.
RDY Priority Override Count ² ³	536	Brocade switch port Brocade Inter-switch connection ¹ Brocade switch	Average count of times per second during which the sending of R_RDY or VC_RDY signals was a higher priority than the sending of frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. When you troubleshoot a SAN, use this count to view port conditions that might slow the performance of the resources to which those ports are connected.
Port State Change Rate	537	Brocade switch port Brocade Inter-switch connection ¹ Brocade switch	Average count of times per second that the state of a port changes to offline, online, or faulty. When you troubleshoot a SAN, use this count to view port conditions that might slow the performance of the resources to which those ports are connected.
Bad EOF CRC Error Count ²	538	Fabrics	The number of cyclic redundancy check (CRC) errors that were detected in frames with a bad end-of-frame (EOF) indicator. In Brocade fabrics, a bad EOF in a frame can indicate that the frame has a known, previously detected CRC error. A good EOF in a frame with a CRC error indicates that the CRC error was not previously detected.
Port Abort Count	539	DS8000: port	The number of times that the port on DS8000 received an abort error.
Extreme I/O Concurrency Count	540	DS8000: port	The number of times that the port on DS8000 had more than 1500 concurrent I/O operations or exchanges. The number of concurrent I/O operations for a port on DS8000 cannot exceed 2000.
I/O Busy Count	541	DS8000: port	The number of times that the port on DS8000 returned a SCSI Queue Full or a Busy status to the server. Ports can return these statuses if the number of concurrent I/O operations or exchanges exceeds an internal DS8000 threshold.
I/O Overrun Count	542	DS8000: port	The number of times that the port on DS8000 had to discard commands because the number of concurrent I/O operations or exchanges for the port exceeded 2000.
Zero Send Buffer Credit Time ²	543	DS8000: port	The number of one-second intervals during which the port on DS8000 had depleted its send buffer credits. That is, the approximate amount of time that the receiving port had no credit to provide to the port on DS8000.
Zero Receive Buffer Credit Time ²	544	DS8000: port	The number of one-second intervals during which the port on DS8000 had depleted its receive buffer credits. That is, the approximate amount of time that the port on DS8000 had no credit to provide to the sending port.
Port Transmit Delay Time	547	IBM Spectrum Virtualize: port	The duration of the delay on the port during send operations, in microseconds. The reason for the delay might be a lack of buffer credits.
Port Transmit Delay Count	548	IBM Spectrum Virtualize: port	The number of send operations by the port that were delayed. The reason for these delays might be a lack of buffer credits.
Port Transmit Delay Measured Count	549	IBM Spectrum Virtualize: port	The number of send operations by the port that were measured to allow the detection of any potential transmission delay due to lack of buffer credits.
<p>Note:</p> <ol style="list-style-type: none"> Performance data for inter-switch connections contains performance data for the following resources: <ul style="list-style-type: none"> Switch ports on ISLs and NPV links. Switch trunks for ISL trunks, ICL trunks, F_port channels, and port channels. The performance data for both resources is displayed together in the IBM Spectrum Control GUI. However, to obtain this data from the CLI, you must run separate queries. This counter is only available for ports on Brocade switches. You must use an SNMP agent to collect data for this counter. 			

Command-line interface

The following sections describe the IBM Spectrum® Control command-line interface (CLI).

The following topics are included.

- CLI requirements
- Command modes
- Syntax diagram conventions that are used in this guide
- Tptool commands
- Command aliases
- Parameter aliases

- [CLI requirements](#)

Verify the requirements for running the command-line interpreter and for using the command-line interface.

- [Command modes](#)
You can use the CLI to run one command or a series of commands, either interactively or from a script.
- [Conventions used in this guide](#)
Information is given about the conventions that are used in this publication.
- [tpctool command](#)
The **tpctool** command is the IBM Spectrum Control CLI program. The IBM Spectrum Control **tpctool** command can be used either on its own, using the associated options and arguments, or interactively by starting the **tpctool** command with no options or arguments to start an interactive session.
- [Command aliases](#)
This topic discusses command aliasing.
- [Parameter aliases](#)
This topic lists common parameters and their aliases.

CLI requirements

Verify the requirements for running the command-line interpreter and for using the command-line interface.

- Verify that the installer checks for and installs the correct version of Java™, and configures the CLI to use the installed Java.
- Verify that IBM Spectrum® Control is installed and running, and that storage devices are discovered.
- Verify that the CLI is connected to a IBM Spectrum Control Device server before you use all commands. The **encrypt** command is the only exception that does not require the connection.
- Verify that you have a valid user ID, password, and URL before you use all commands, except for the **encrypt** command.

Command modes

You can use the CLI to run one command or a series of commands, either interactively or from a script.

Single-shot mode

To run a single command, specify the CLI program and that command at the shell prompt, for example:

```
shell> tpctool lsdev -user me -pwd mypass -url myhost:myport -l -sys
```

Interactive mode

To run in interactive mode, enter the **tpctool** command with no command-line options. In the following example, the string "shell>" represents the shell prompt and "tpctool>" represents the prompt from the CLI while in interactive mode. At the CLI prompt, any valid CLI command can be entered.

```
shell> tpctool
tpctool> lsdev -user me -pwd mypass -url myhost:myport -l -sys
```

Script mode

To run a set of commands that you defined in a file, start the CLI program and specify a file that contains the commands, for example:

```
shell> tpctool -script ~/bin/containersetup
shell>
```

Conventions used in this guide

Information is given about the conventions that are used in this publication.

This publication uses several conventions for special terms and actions, and for operating system-dependent commands and paths. The following typeface conventions are used in this publication:

Bold

- Flags that display with text
- Graphical user interface (GUI) elements (except for titles of windows and dialogs)
- Names of keys

Italic

- Variables
- Values that you must provide
- New terms
- Words and phrases that are emphasized
- Titles of documents

monospace

- Commands and command options
- Flags that display on a separate line
- Code examples and output
- Message text
- Names of files and directories
- Text strings that you must type, when they display within text
- Names of Oracle Java™ methods and classes
- HTML and XML tags that display **like this**, in monospace type

For syntax notations, remember the following details.

- In AIX®, the prompt for the root user is #.
- In AIX and Linux®, the commands are case-sensitive, so you must type commands exactly as they are shown.
- **Syntax diagram conventions**
A syntax diagram uses symbols to represent the elements of a command and to specify the rules for using these elements.

Syntax diagram conventions

A syntax diagram uses symbols to represent the elements of a command and to specify the rules for using these elements.

Syntax diagrams use position to indicate required, optional, and default values for keywords, variables, and operands.

A keyword represents the name of a command, flag, parameter, or argument. Required keywords indicate the parameters or arguments that must be specified for the command.

To read syntax diagrams, follow the path of the line. Read the diagrams from left-to-right, top-to-bottom, following the main path line.

Main path line

The main path line begins on the left with double arrowheads (>>) and ends on the right with two arrowheads facing each other (<>). If a diagram is longer than one line, each line to be continued ends with a single arrowhead (>) and the next line begins with a single arrowhead. The -->< symbol indicates the end of the syntax diagram.


Required keywords

Required keywords appear on the main path line. Mutually exclusive required keywords are stacked vertically. In the following example, you must choose A, B, and C.


➤ A — B — C ➤

Optional keywords

Optional keywords indicate parameters or arguments that you might choose to specify for the command. Optional keywords appear below the main path line. Mutually exclusive optional keywords are stacked vertically. In the following example, you must choose A, B, or C.



When an optional item appears above the main line, the item above the line is the default value when no optional item is specified in the command. In the following example, the user has the same choices as above (A, B, C, or nothing at all), but if nothing is selected, the default value will be A.




Repeatable items

A stack of items followed by an arrow returning to the left means that you can select more than one item or, in some cases, repeat a single item. For example:


An arrow returning to the left means you can repeat the item.



If one or more characters appear in the arrow's line, those characters are required as a delimiter to separate repeated items.



If you can choose from two or more items, they are displayed vertically in a stack. A stack of items followed by an arrow returning to the left means that you can select more than one item or, in some cases, repeat a single item. In the following example, you can choose any combination of A, B, or C.



Variables

Italicized, lowercase elements denote variables. In the following example, you must specify a variable value when you enter the keyword command:

➤ keyword — *variable* ➤

tpctool command

The **tpctool** command is the IBM Spectrum Control CLI program. The IBM Spectrum Control **tpctool** command can be used either on its own, using the associated options and arguments, or interactively by starting the **tpctool** command with no options or arguments to start an interactive session.

Syntax

➡ **tpctool** — *command* — *connection-options* — *formatting-options* — *command-options* — *command-arguments* →

Parameters

command

Either one of the commands (generic, disk, fabric, or reporting) or a command alias.

connection-options

Options that are used to connect to IBM Spectrum® Control. These options are included.

- -user
- -pwd
- -url

These options are described with details in the **tpctool** arguments and options list.

formatting-options

Options that are used to format output for some of the commands.

command-options

Options that are associated with commands that define command behavior.

command-arguments

Generally, values used to define command options. They can be defined as lists of string values that are separated by commas or spaces.

The following arguments and options are valid for **tpctool**:

-user *user_name*

Specifies an IBM Spectrum Control user ID.

-pwd *password*

Specifies the password for the IBM Spectrum Control user ID.

-url *url*

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-help | -h | -?

Lists help information for the command.

-ver

Displays the version of the installed IBM Spectrum Control.

Command aliasing

The IBM Spectrum Control CLI provides the capability for command aliasing by using the command configuration file. The default configuration file is `c:\program files\ibm\tpc\cli\libs\tpccli.conf`. With aliasing, you define a name for the alias followed by a value that is the name of a command and any options that are associated with command. The aliased command string is replaced by the defined value and the entire line is parsed again. Passwords that are used in aliased commands must first be encrypted by using the **encrypt** command.

For example, to shorten a frequently used command that you can define the following alias:

```
tpctool>lsperf = lsdev -user dsadmin -pwd dsadmpw1 -url 9.44.33.126:8990 -fabric -perf
```

After the alias is defined, you can run the **lsperf** command to run the aliased **lsdev** command.

You can provide a short form command that targets different Device servers, as follows:

```
tpctool>lsperf1 = lsdev -user dsadmin -pwd dsadmpw1 -url hostOne:9161 -perf
```

```
tpctool>lsperf2 = lsdev -user dsadmin -pwd dsadmpw1 -url hostTwo:9161 -perf
```

You can specify more options and arguments for an aliased command.

```
tpctool>lsperf2 -fabric -ctype port
```

Which expands to:

```
tpctool>lsdev -user dsadmin -pwd dsadmpw1 -url hostTwo:9161 -perf -fabric -ctype port
```

You can also nest aliases:

```
tpctool>lsperf = lsdev -user dsadmin -pwd dsadmpw1 -url hostTwo:9161 -perf -fabric
```

```
tpctool>lsperf1 = lsdev -user dsadmin -pwd dsadmpw1 -url hostOne:9161
```

```
tpctool>lsperf2 = lsdev -user dsadmin -pwd dsadmpw1 -url hostTwo:9161
```

To unset an alias, type the name of the command alias followed by the equals (=) sign.

```
lsperf =
```

When you use an alias with a key and value pair in the `tpccli.conf` file, you must specify four back slashes for each back slash.

For example, this value:

```
ABCDE1 = CLARiON\+ABC01234567890+0
```

Is specified as:

Examples

To start an interactive session for the IBM Spectrum Control CLI:

The following command starts an interactive session. After the session starts, you can use other CLI commands in the session:

```
tpctool
```

To start an interactive session with credentials:

The following command starts an interactive session and specifies the user ID, password, and destination:

```
tpctool -user adminpass
-pwd adminpw1
-url 9.43.124.255:8080
```

Setting up credentials for automatic log in

You can set up and save your credentials in a different CLI program configuration file other than the default file without entering the interactive mode to eliminate the requirement of specifying the credentials again for every **tpctool** command session.

Creating a new configuration file

You can set a new location and file name and create a new configuration file that contains your credentials for logging in to the **tpctool** command.

This is the default configuration file location for the Windows operating system:

```
C:\program files\ibm\tpc\cli\libs\tpccli.conf
```

For the Linux® operating system:

```
/opt/IBM/TPC/cli/libs/tpccli.conf
```

Perform these steps to create a new configuration file:

1. In a command prompt window, enter the following command, for a Windows operating system:

```
SET TPCCFG=C:\installation_dir\cli\config file
```

Where the `cli.properties` file is the default configuration file that already exists and `TPCCFG` is the new name of the configuration file.

For a Linux operating system:

```
export TPCCFG=../installation_dir/cli/config file
```

Where the `cli.properties` file is the default configuration file that already exists and `TPCCFG` is the new name of the configuration file.

2. Edit your new configuration file to include a destination, your user ID and password.

For example, a Windows operating system:

```
C:\installation_dir\cli>tpctool
tpctool> url=localhost:9550
tpctool> user=db2admin
tpctool> password=xxxxxxxxx
```

Where:

```
url=localhost:9550
```

The name or IP address of your local or remote server.

```
user=db2admin
```

The DB2® instance name that is used when installing IBM Spectrum Control.

```
password=xxxxxxxxx
```

The password that you use to log in to IBM Spectrum Control.

For example, a Linux operating system:

```
</installation_dir/cli/tpctool.sh
tpctool> url=localhost:9550
tpctool> user=db2inst1
tpctool> password=xxxxxxxxx
```

3. Exit the **tpctool** session:

```
tpctool> quit
```

When you exit the command line interface on a Linux operating system, and, then want to access the command again, you need to run the **export** command again.

For example:

```
export TPCCFG=../installation_dir/cli/config file
```

You can interactively create a new configuration file. Once you locate the IBM Spectrum Control installation directory, double-click the **tpctool.bat** file or enter **tpctool.bat|sh file** on a command prompt window to start the **tpctool** session in the interactive mode. Enter your credentials to alleviate the need to specify the credentials again, when you start each **tpctool** command.

For example, on a Windows system, the **tpctool.bat** file is located here:

```
c:\program files\ibm\tpc\cli directory file
```

On Linux operating systems, you must add the directory path as `/opt/IBM?TPC/c11` to the `tpctool.sh` file. The command input can be either lowercase, uppercase, or mixed case, unless specific command parameters require case sensitivity.

Return codes

The following table contains the codes that are returned by the **tpctool** command.

Table 1. Return codes for the **tpctool** command

Code	Description
0	The command completed successfully.
1	The command was unknown to tpctool and was not resolved as an alias.
2	A required option was not provided.
3	An option was unknown to tpctool or was not applicable to the command.
4	An option was missing a required parameter.
5	The format of a parameter for the option was not valid.
6	The format of an argument was not valid.
7	An extraneous argument or argument list was provided.
8	The tpctool client could not connect with the Device server.
9	The tpctool client could not log in to tIBM Spectrum Control using the specified credentials.
10	The specified credentials are not authorized to perform the requested action.
11	A required component is not installed and enabled.
12	The command might have started, but the connection with IBM Spectrum Control was lost. The command might not be completed successfully.
13	Some operations were partially completed before IBM Spectrum Control returned a failure.
14	The command failed.
15	An attempt was made to remove a group that has a child group associated with it. Remove the child group and try the operation again.

- [actzs](#)
Use the **actzs** command to activate changes to the zone set in the fabric. This command must be run within a transaction. You must have Administrator authority to use this command.
- [addza](#)
Use the **addza** command to add a zone alias to a zone. You must have Administrator authority to use this command.
- [addzaptops](#)
Use the **addzaptops** command to add ports to a zone alias. You must have Administrator authority to use this command.
- [addzone](#)
Use the **addzone** command to add a zone to a zone set. This command must be run as a transaction. For more information, see the **start** command. You must have Administrator authority to use this command.
- [addzoneports](#)
Use the **addzoneports** command to add switch ports to a zone. This command must be run as a transaction. You must have Administrator authority to use this command.
- [assignvol](#)
Use the **assignvol** command to assign host ports to volumes. You must have Administrator authority to use this command.
- [autosetarray](#)
Use the **autosetarray** command to extract and save information about storage pools.
- [catdscfg](#)
Use the **catdscfg** command to list the contents of the property files for the Device server and to check the status of the Device server. You must have Administrator authority to use this command.
- [chexport](#)
Use the **chexport** command to change the protocol configuration for an export.
- [chfs](#)
Use the **chfs** command to change the properties of a file system, and to add or remove disks to or from a file system.
- [chfset](#)
Use the **chfset** command to change the attributes of a fileset.
- [chkquota](#)
Use the **chkquota** command to check quota limits for users, user groups, and filesets on a file system and to write the information to the GPFS database.
- [chwcach](#)
Use the **chwcach** command to modify the attributes of a cache fileset on a wide area network (WAN) cluster.
- [chwcachesource](#)
Use the **chwcachesource** command to change the properties of a home system on a fileset. The home system is the source of the data in a wide area network (WAN) cache configuration. You can use WAN caching to distribute data transparently among data centers and multiple remote locations without disruption to applications. You must have Administrator authority to use this command.
- [ckzone](#)
Use the **ckzone** command to verify that a fabric contains a zone.
- [ckzs](#)
Use the **ckzs** command to verify that a fabric contains a zone set.
- [commit](#)
Use the **commit** command to commit a transaction. When you commit a transaction, all the commands issued after you started the transaction are enacted. You must have Administrator authority to use this command.
- [ctlwcache](#)
Use the **ctlwcache** command to run maintenance operations on a cache fileset on a wide area network (WAN) cluster. Data on a source cluster is cached to this fileset. You can resynchronize cache data to the source fileset, flush the cache to the source fileset, remove cached data, and expire or unexpire the data in a fileset. Use the **lswcache** command to view the attributes of the cache fileset. You must have Administrator authority to use this command.
- [deactzs](#)
Use the **deactzs** command to deactivate the active zone set. This command must be run as a transaction. You must have Administrator authority to use this command.
- [encrypt](#)
Use the **encrypt** command to generate an encrypted password for use in the configuration file. This command takes text from standard input and generates 7-bit ASCII-equivalent characters (uuencode).

- [**getdscfg**](#)
Use the **getdscfg** command to list the current value of a property from the property file for the Device server. You must have Administrator authority to use this command.
- [**getrpt**](#)
Use the **getrpt** command to list a performance report for a specified storage subsystem.
- [**linkset**](#)
Use the **linkset** command to create a junction to connect a name in a directory of a parent file set to the root directory of a child fileset. You must have Administrator authority to use this command.
- [**lsappgroup**](#)
Use the **lsappgroup** command to display a list of applications that are known to IBM Spectrum Control.
- [**lsappgroupmembers**](#)
Use the **lsappgroupmembers** command to list members of a specified application.
- [**lsarray**](#)
Use the **lsarray** command to display information about arrays and back-end storage systems. For SAN Volume Controller, XIV® systems, and Dell EMC VMAX and VNX storage systems, the command displays information about block pools.
- [**lsbackenddisktypes**](#)
Use the **lsbackenddisktypes** command to list the types of back-end disks and their average input/output.
- [**lsbackendraidtypes**](#)
Use the **lsbackendraidtypes** command to list the types of back-end RAID arrays that are available for managed disk groups.
- [**lsbackendtypes**](#)
Use the **lsbackendtypes** command to list the types of back-end storage systems.
- [**lscluster**](#)
Use the **lscluster** command to list all clusters or specified clusters that are on a Storwize® V7000 Unified storage system.
- [**lscomp**](#)
Use the **lscomp** command to list the components for which performance data is collected.
- [**lscounters**](#)
Use the **lscounters** command to list available performance counters.
- [**lsdeptgroup**](#)
Use the **lsdeptgroup** command to display a list of departments that are known to IBM Spectrum Control.
- [**lsdeptgroupmembers**](#)
Use the **lsdeptgroupmembers** command to list members of a specified department.
- [**lsdev**](#)
Use the **lsdev** command to list information about storage systems, fabrics, and switches. This information includes the globally unique identifier (GUID) or worldwide name (WWN) for the fabric, the user-defined name, the resource type, the status, and the time that the status was updated.
- [**lsdevp**](#)
Use the **lsdevp** command to list worldwide port names (WWPNs) for a subsystem.
- [**lsdisk**](#)
Use the **lsdisk** command to list all the physical disks on a specified storage subsystem.
- [**lsexport**](#)
Use the **lsexport** command to list all exports or specified exports that are associated with a cluster that is on a Storwize V7000 Unified storage system. You can also list exports by file system, Netapp, or server that is managed by a Storage Resource Agent (SRA).
- [**lsextent**](#)
Use the **lsextent** command to display a list of all the storage extents on a specified storage system. An example of a storage extent is a managed disk on a SAN Volume Controller. For Dell EMC VMAX storage systems, the command displays the Disk Group.
- [**lsfcpath**](#)
Use the **lsfcpath** command to list the paths for data transmission between a system with a fibre-channel host bus adapter (HBA) and a storage subsystem.
- [**lsfs**](#)
Use the **lsfs** command to list all file systems or specified file systems that are associated with a cluster on a Storwize V7000 Unified or IBM Spectrum Scale storage system. You can also list file systems by pool or Network Shared Disk (NSD).
- [**lsfset**](#)
Use the **lsfset** command to list all filesets or specified filesets that are associated with a file system on a Storwize V7000 Unified or IBM Spectrum Scale storage system.
- [**lshtype**](#)
Use the **lshtype** command to list host types. You must have Administrator authority to use this command.
- [**lsoptschedules**](#)
Use the **lsoptschedules** command to show a list of the schedules that you created to analyze storage tiering.
- [**lsmetrics**](#)
Use the **lsmetrics** command to list available performance metrics. You must have Fabric operator or Disk operator authority to use this command.
- [**lsnode**](#)
Use the **lsnode** command to list all nodes or specified nodes that are associated with a cluster that is on a Storwize V7000 Unified storage system.
- [**lsnsd**](#)
Use the **lsnsd** command to list all Network Shared Disks (NSDs) or specified NSDs that are on a Storwize V7000 Unified storage system. You can also list NSDs by file system, pool, or both.
- [**lspool**](#)
Use the **lspool** command to list all file system pools that are on a specified Storwize V7000 Unified storage system. You can also list the pools by file system.
- [**lsport**](#)
Use the **lsport** command to list the ports that are on a Fibre Channel host bus adapter (HBA).
- [**lsquota**](#)
Use the **lsquota** command to list all quotas or specified quotas that are on a Storwize V7000 Unified storage system. You can also list quotas by file system.
- [**lssrg**](#)
Use the **lssrg** command to display a list of storage resource groups that are known to IBM Spectrum Control.
- [**lssrgmembers**](#)
Use the **lssrgmembers** command to list members of a specified storage resource group.
- [**lssvr**](#)
Use the **lssvr** command to list all systems that are discovered by Fabric Manager. You must have Fabric Administrator authority to use this command.
- [**lssvrdisk**](#)
Use the **lssvrdisk** command to list all of the physical disks that are known to a specified server.
- [**lsswitch**](#)
Use the **lsswitch** command to display a list of all switches in a specified fabric.

- [**lstime**](#)
Use the **lstime** command to list the time ranges for which performance data is available.
- [**lstype**](#)
Use the **lstype** command to list the components that are recognized by IBM Spectrum Control. No authorization is required to run this command.
- [**lsvm**](#)
Use the **lsvm** command to list all virtual machines that are known to IBM Spectrum Control. You can also list only the virtual machines that were discovered through hypervisors.
- [**lsvmdisk**](#)
Use the **lsvmdisk** command to list all the VMWare virtual machine disks that are known to the given virtual machine or hypervisor.
- [**lsvol**](#)
Use the **lsvol** command to list all volumes on a system, list a specific volume or volumes, or list volumes on a specific array.
- [**lsvolgroup**](#)
Use the **lsvolgroup** command to display a list of volume groups that are known to IBM Spectrum Control.
- [**lszone**](#)
Use the **lszone** command to list the zones in a zone set.
- [**lszs**](#)
Use the **lszs** command to list information about zone sets. This information includes the zone set name and status.
- [**mkappgroup**](#)
Use the **mkappgroup** command to create an application with a specific name, type, description, and user-defined properties. You must have Administrator authority to use this command.
- [**mkdeptgroup**](#)
Use the **mkdeptgroup** command to create a department with a specific name, description, type, subtype, and user-defined properties. You must have Administrator authority to use this command.
- [**mkexport**](#)
Use the **mkexport** command to create an export to access data through a data transfer protocol.
- [**mkfs**](#)
Use the **mkfs** command to create a GPFS file system to manage files on a storage device.
- [**mkfset**](#)
Use the **mkfset** command to create a fileset that is associated with a file system on a Storwize V7000 File Module storage system. With filesets, you can use functions such as snapshots or quotas within a file system. You must have Administrator authority to use this command.
- [**mksrg**](#)
Use the **mksrg** command to create a storage resource group with a specific name, type, description, and user-defined properties. You must have Administrator authority to use this command.
- [**mkzone**](#)
Use the **mkzone** command to create a zone. This command must be run within a transaction. You must have Administrator authority to use this command.
- [**mkzs**](#)
Use the **mkzs** command to create a zone set. This command must be run as a transaction. You must have Administrator authority to use this command.
- [**modifyappgroup**](#)
Use the **modifyappgroup** command to modify an existing application. You can add and remove members of the application by specifying the member type and the key for the individual member. You can also use a -tagkey and -tagvalue member pair.
- [**modifyappgroupviafile**](#)
Use the **modifyappgroupviafile** command to modify an existing application by using an input file. You can add and remove members by specifying the member type and identifiers or by using a member type with a -tagkey and -tagvalue pair.
- [**modifydeptgroup**](#)
Use the **modifydeptgroup** command to modify an existing department. You can add and remove members of a department by specifying the member type. For example, a deptgroup type for a department or an appgroup type for an application and the -memberid parameter for the individual member. You can also use a -tagkey and -tagvalue member pair.
- [**modifydeptgroupviafile**](#)
Use the **modifydeptgroupviafile** command to modify an existing department by using an input file. You can add and remove members by specifying the member type (deptgroup, appgroup) and identifiers or by using a member type with a -tagkey and -tagvalue pair.
- [**modifysrg**](#)
Use the **modifysrg** command to modify an existing storage resource group. You can add and remove members by specifying the member type, such as switch or volume, and the key for the member you want to add or remove.
- [**mountfs**](#)
Use the **mountfs** command to mount a file system on all interface and management nodes or a specified subset. You must have Data Administrator authority to use this command.
- [**rmappgroup**](#)
Use the **rmappgroup** command to delete a specific application and its members. If you use the **rmappgroup** command without the -rmchildren parameter, the application is deleted, and the members are moved up a level in the hierarchy.
- [**rmbackenddisktype**](#)
Use the **rmbackenddisktype** command to remove a back-end type of disk.
- [**rmbackendraidtype**](#)
Use the **rmbackendraidtype** command to remove a back-end RAID array type from the set of RAID array types available for managed disk groups.
- [**rmbackendtype**](#)
Use the **rmbackendtype** command to remove a type of back-end storage system.
- [**rmdeptgroup**](#)
Use the **rmdeptgroup** command to delete a specific department and its members. If you use the **rmdeptgroup** command without the -rmchildren parameter, the department is deleted, and the members are moved up a level in the hierarchy.
- [**rmexport**](#)
Use the **rmexport** command to remove an export.
- [**rmfs**](#)
Use the **rmfs** command to remove a file system from an active management node. You must have Administrator authority to use this command.
- [**rmfset**](#)
Use the **rmfset** command to remove a fileset from a file system. You must have Administrator authority to use this command.
- [**rmsrg**](#)
Use the **rmsrg** command to delete the specified Storage Resource Group. The group types (Storage Resource Groups, Reporting Groups and Application Groups) that have child groups are not deleted. A warning message is issued and indicates that the group type being deleted contains child groups and the child groups must be deleted before the parent group can be deleted.
- [**rmza**](#)
Use the **rmza** command to remove a zone alias or aliases from a zone. You must have Administrator authority to use this command.

- [rmzaptops](#)
Use the **rmzaptops** command to remove a port or ports from a zone alias. You must have Administrator authority to use this command.
- [rmzone](#)
Use the **rmzone** command to delete a zone or remove a zone from a zone set. If you remove or delete the last zone in a zone set, the zone set is also deleted. This command must be run as a transaction. You must have Administrator authority to use this command.
- [rmzoneports](#)
Use the **rmzoneports** command to remove switch ports from a zone. This command must be run as a transaction. You must have Administrator authority to use this command.
- [rmzs](#)
Use the **rmzs** command to delete a zone set. This command must be run as a transaction. You must have Administrator authority to use this command.
- [rollback](#)
Use the **rollback** command to erase any commands that were issued since you started the transaction. You must have Administrator authority to use this command.
- [runoptschedule](#)
Use the **runoptschedule** command to run a schedule that analyzes storage tiering.
- [setarray](#)
Use the **setarray** command to set the type of back-end storage system, type of Redundant Array of Independent Disks (RAID), type of disk, and number of disks for an array.
- [setbackenddisktype](#)
Use the **setbackenddisktype** command to set or update the type of back-end disk.
- [setbackendraidthe](#)
Use the **setbackendraidthe** command to set the types of back-end RAID arrays for managed disk groups.
- [setbackendtype](#)
Use the **setbackendtype** command to set or update the type of back-end storage system.
- [setdscfg](#)
Use the **setdscfg** command to set the value of a property in the property file for the Device server. You must have Administrator authority to use this command.
- [setquota](#)
Use the **setquota** command to set the amount of disk space and number of inodes that are assigned on a file system for a specified user name, group, or fileset.
- [showoptresults](#)
Use the **showoptresults** command to show the recommendations for optimizing the placement of volumes.
- [start](#)
Use the **start** command to start a transaction. You must have Administrator authority to use this command.
- [unassignvol](#)
Use the **unassignvol** command to remove the host ports from the assignment list for a volume. You must have Administrator authority to use this command.
- [unlinkset](#)
Use the **unlinkset** command to unlink a fileset. You must have Administrator authority to use this command.
- [unmountfs](#)
Use the **unmountfs** command to unmount a file system.
- [updatesrg](#)
Use the **updatesrg** command to update the attributes of the specified Storage Resource Group with a new name, description, or user-defined properties.

Related information

- [Parameter aliases](#)
- [Command aliases](#)

actzs

Use the **actzs** command to activate changes to the zone set in the fabric. This command must be run within a transaction. You must have Administrator authority to use this command.

Syntax

```

▶▶ tpc tool — actzs — -user user_name -pwd password -url url — -fabric — WWN —
                                     -help — ? — -silent —▶
▶ — zone_set —▶▶

```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
 - pwd password
Specifies the password for the IBM Spectrum Control user ID.
 - url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
 - fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
 - help | -h | -?
Lists help information for the command.
 - silent
Suppresses all output for the command. If you omit this parameter, output is included.
- zone_set

The `zone_set` variable is the name of the zone_set.

Example: Activating changes to the zone set

The following commands activate the PARIS zone set:

```
tpctool> -user me -pwd mypass -url myhost:myport
tpctool> start -fabric 100000051E34F6A8
tpctool> actzs -fabric 100000051E34F6A8 PARIS
tpctool> commit -fabric 100000051E34F6A8
```

addza

Use the **addza** command to add a zone alias to a zone. You must have Administrator authority to use this command.

Syntax

```
► tpctool — addza — -user user_name -pwd password -url url — -fabric WWN — -help — ? — -silent — -zone →
► zone — alias ◀
```

Parameters and arguments

`-user user_name`
Specifies an IBM Spectrum Control user ID.

`-pwd password`
Specifies the password for the IBM Spectrum Control user ID.

`-url url`
Specifies the Device server. The format of the URL is `system:port_number`, where `system` represents either the host name or IP address, and `port_number` represents the IBM Spectrum Control Device server port.

`-fabric WWN`
Specifies the fabric. The `WWN` variable is the worldwide name (WWN).

`-help | -h | -?`
Lists help information for the command.

`-silent`
Suppresses all output for the command. If you omit this parameter, output is included.

`-zone zone`
Specifies the name of the zone where you want to add a zone alias.

`alias`
Specifies the name of the zone alias to be added to the zone.

Example: Adding a zone alias

The following commands add the PARIS zone alias to the EUROPE zone:

```
tpctool -user me -pwd mypass -url myhost:myport
tpctool> start -fabric 100000051E34F6A8
tpctool> addza -fabric 100000051E34F6A8 -zone EUROPE PARIS
tpctool> commit -fabric 100000051E34F6A8
```

addzaptops

Use the **addzaptops** command to add ports to a zone alias. You must have Administrator authority to use this command.

Syntax

```
► tpctool — addzaptops — -user user_name -pwd password -url url — -fabric WWN — -help — ? — -silent — -za — zone_alias — port ◀
```

Parameters and arguments

`-user user_name`
Specifies an IBM Spectrum Control user ID.

`-pwd password`
Specifies the password for the IBM Spectrum Control user ID.

`-url url`
Specifies the Device server. The format of the URL is `system:port_number`, where `system` represents either the host name or IP address, and `port_number` represents the IBM Spectrum Control Device server port.

-fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).

-help | -h | -?
Lists help information for the command.

-silent
Suppresses all output for the command. If you omit this parameter, output is included.

-za *zone_alias*
Specifies the name of the zone alias where the port is to be added.

port
Specifies the name of the port to be added to the zone alias.

Example: Adding a port to a zone alias

The following commands add a port to the PARIS zone alias:

```
tpctool -user me -pwd mypass -url myhost:myport
tpctool> start -fabric 100000051E34F6A8
tpctool> addzports -fabric 100000051E34F6A8 -za PARIS 210000E08B0B4C2G
tpctool> commit -fabric 100000051E34F6A8
```

addzone

Use the **addzone** command to add a zone to a zone set. This command must be run as a transaction. For more information, see the **start** command. You must have Administrator authority to use this command.

Syntax

```
tpctool — addzone — -user user_name -pwd password -url url — -fabric — WWN — -help — ? — -silent — -zs — zone_set — zone —
```

Parameters and arguments

-user *user_name*
Specifies an IBM Spectrum Control user ID.

-pwd *password*
Specifies the password for the IBM Spectrum Control user ID.

-url *url*
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).

-help | -h | -?
Lists help information for the command.

-silent
Suppresses all output for the command. If you omit this parameter, output is included.

-zs *zone_set*
Specifies the zone set. The *zone_set* variable is the name of the zone set.

zone
Specifies the zone.

Example: Add a zone to a zone set

The following commands add the WINDOWSNT zone to the PARIS zone set.

```
tpctool> start -user me -pwd mypass -url myhost:myport -fabric 100000051E34F6A8
tpctool> addzone -fabric 100000051E34F6A8 -zs PARIS WINDOWSNT
tpctool> commit -fabric 100000051E34F6A8
```


Related reference

- [start](#)

addzoneports

Use the **addzoneports** command to add switch ports to a zone. This command must be run as a transaction. You must have Administrator authority to use this command.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- zone zone
Specifies the zone. The *zone* variable is the name of the zone.
- help | -h | -?
Lists help information for the command.
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- ports | -
Specifies the switch ports. The *ports* variable is a list of worldwide port names (WWPNs). If you specify a single dash (-), the WWPNs are read from standard input.

Example: Adding a switch port to a zone set


The following commands add several switch ports to the WINDOWSNT zone. The list of WWPNs is read from standard input.

```
tpctool> start -user me -pwd mypass -url myhost:myport -fabric 100000051E34F6A8
tpctool> addzoneports -fabric 100000051E34F6A8 -zone WINDOWSNT -
tpctool> commit -fabric 100000051E34F6A8
```

assignvol

Use the **assignvol** command to assign host ports to volumes. You must have Administrator authority to use this command.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- dev subsystem
Specifies the globally unique identifier (GUID) of a storage subsystem such as that obtained by running the **lsdev -subsys** command.
- hp host_port
Specifies the host ports. The *host_port* variable is a comma-separated list of worldwide port numbers (WWPNs), such as the port numbers obtained by running the **lsport** command.
- hn host_name
Specifies the names of the host systems for the ports. The *host_name* variable is a comma-separated list of host system names. There must be one name for each port in the list of host ports.
- ht host_type
Specifies the type of host system on which the Fibre Channel port is located. The *host_type* variable is a comma-separated list of host system types, such as the types obtained by running the **lshtype** command.
- lun LUN_ID
Specifies the logical unit number (LUN ID) that the hosts map to the volumes. The *LUN_ID* variable is a comma-separated list of LUN IDs. There must be one LUN ID for each volume in the argument list.
- sp subsystem_port

- Specifies the worldwide port numbers (WWPNs) that the hosts use to access the volume. The *subsystem_port* variable is a comma-separated list of WWPNs, such as that obtained by running the **lsdevp** command. If you omit this parameter, the default ports are used.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- volume id | -
Specifies the volumes. The *volume_ID* variable is a comma-separated list of volume IDs, such as that obtained by running the **lsvol** command. If a single dash (-) is issued, the volume IDs are read from standard input.

Example: Assign a host port to a volume

The following command assigns a host port to a volume:

```
tpctool assignvol -user me -pwd mypass -url myhost:myport -dev 2105.22232+0
-hp 5005076300C79470 -lun a3
```

Related reference

- [lsdev](#)
- [lsport](#)
- [lshtype](#)
- [lsdevp](#)

autosetarray

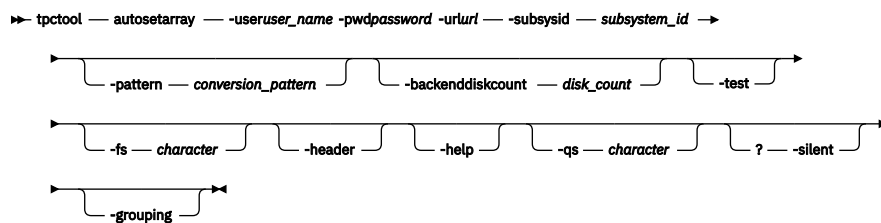
Use the **autosetarray** command to extract and save information about storage pools.

You can save information about a storage pool such as the type of back-end storage system, type of RAID, and type of disk. To do this action, use a default pattern or create a pattern that matches the names of pools in a subsystem. If the match is successful, the information about the pool is extracted and saved. This command is available for the following storage systems.

- Storwize® V7000
- Storwize V7000 Unified
- SAN Volume Controller

You must have Administrator authority to use this command.

Syntax



Storage pool names and pattern matching

The names of storage pools must contain the following information:

- A character that represents a valid type of back-end storage system
- A number or character that represents a valid type of back-end RAID
- A series of characters that represent a valid type of back-end disk

Tip: To obtain a list of valid types of back-end storage system, types of disk, and types of RAID, run the following commands: **lsbackendtypes**, **lsbackenddisktypes**, and **lsbackendraitypes**. If the types of back-end storage system, types of disk, and types of RAID that are used in the pool names are not listed, you can add them to the list. See the descriptions of the following commands: **setbackentype**, **setbackenddisktype**, and **setbackendraitype**.

The default pattern is structured as shown in the following table.

Position in pattern	Pattern
1	B
2	x
3	x
4	x

Position in pattern	Pattern
5	x
6	x
7	x
8	D
9	R
10	D
11	D
12	x*

You can use the default pattern or create patterns to match the names of pools in a storage subsystem. A list of the characters that are used to create patterns is provided in the following table.

Character	Represents
B	The character that is used to represent the type of back-end storage system. The matching character must meet the following criteria. <ul style="list-style-type: none"> • Must be in the position that is specified by the pattern • Must be of the same length • Must be a valid type of storage system In the default pattern, one character is used to represent the type of back-end storage system.
x	The character that is used to represent insignificant characters. In the default pattern, the second, third, fourth, fifth, sixth, and seventh characters in a matching pool name are ignored.
D	The characters that are used to represent the type of back-end disk. Matching characters must meet the following criteria. <ul style="list-style-type: none"> • Must be in the position that is specified by the pattern • Must be of the same length • Must be a valid type of disk In the default pattern, three characters are used to represent the type of back-end disk.
R	The number or character that is used to represent the type of back-end RAID such as 1, or x. The matching character must meet the following criteria. <ul style="list-style-type: none"> • Must be in the position that is specified by the pattern • Must be of the same length • Must be a valid type of RAID In the default pattern, one character is used to represent the type of back-end RAID.
x*	The character followed by an asterisk that is used to represent zero or any number of insignificant characters that occur at the end of a pool name. In the default pattern, the 12th character and all subsequent characters in a matching pool name are ignored.
C	The characters that are used to represent the number of disks. In the default pattern, the number of disks is not specified. To specify the number of disks, use the -backenddiskcount parameter. The default value is 1. You can create a custom pattern to extract the number-of-disks value from the name of the MDisk group. Matching characters must meet the following criteria. <ul style="list-style-type: none"> • Must be in the position that is specified by the pattern • Must use numeric values such as 0 - 9 • Must be of the same length <ul style="list-style-type: none"> ◦ C represents 0 - 9 ◦ CC represents 00 - 99

Sample: Using the default pattern

You issue the **autosetarray** command and the default pattern is used to extract information about the CFG1xGGA1071 storage pool. In the table, each character in the default pattern is matched against the corresponding character in the pool name.

Default pattern	Pool name	Description
B	C	The type of back-end storage system
x	F	The character is ignored
x	G	The character is ignored
x	1	The character is ignored
x	x	The character is ignored
x	G	The character is ignored
x	G	The character is ignored
D	A	The first character of three characters that is used to represent the type of back-end disk
R	1	The type of back-end RAID
D	0	The second character of three characters that is used to represent the type of back-end disk
D	7	The third character of three characters that is used to represent the type of back-end disk
x*	1	This character and all subsequent characters are ignored

The information about the pool is extracted and saved.

Information	Extracted values	Description
Type of back-end storage system	C	EMC Clariion
Type of back-end disk	A07	SATA - 7,500 rpm
Type of back-end RAID	1	RAID 1
Number of disks	1	1

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- subsysid subsystem_id
Specifies the ID or the globally unique identifier of the storage subsystem.
- pattern conversion_pattern
Specifies the pattern that is used to extract information from pool names. If you do not specify a pattern, the default pattern is used.
- backendsdiskcount disk_count
Specifies the number of disks. If you use the default pattern or create a pattern that does not extract the number of disks from pool names, you can enter the number of disks. The default number of disks is 1.
- test
Verifies the pattern that is used with a specified storage subsystem. The information is extracted and displayed, but it is not saved.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- grouping
Enables the grouping of numeric values. For example, in English the value 12000 would display as 12,000. The grouping character is determined by the system locale.

Example: Using patterns to extract and save information about pools

The following command extracts information about pools in the specified storage subsystem that match the default pattern. If a pool name does not match the pattern, an information message is displayed.

```
tpctool> autosetarray -subsysid 00000200A0C0005C+0
```

The following information messages and output are displayed:

```
[AAJ002013E] The pool name mdiskgroup0 is too short to be matched
against the pattern BxxxxxxDRDDx*.
[AAJ002013E] The pool name Cognos is too short to be matched
against the pattern BxxxxxxDRDDx*.
[AAJ002013E] The pool name mdiskgrp1 is too short to be matched
against the pattern BxxxxxxDRDDx*.
[AAJ002013E] The pool name mdiskgrp2 is too short to be matched
against the pattern BxxxxxxDRDDx*.
```

Name of Array	Back-end Type	Back-end RAID Type	Back-end Disk Type
mdiskgroup0	-	-	-
Cognos	-	-	-
mdiskgrp1	-	-	-
mdiskgrp2	-	-	-
Cpool3GA60713	C	6	A07
Cpool2GA607	C	6	A07
Dpool1GA607F	D	6	A07


Back-end	Disk Count	Status
-	No Match	
-	No Match	
-	No Match	
-	No Match	
1	Successful	
1	Successful	
1	Successful	

- [lsbackendtypes](#)
- [setbackendtype](#)
- [lsbackendsdisktypes](#)
- [setbackendsdisktype](#)
- [setbackendraitype](#)
- [lsbackendraitypes](#)

catdscfg

Use the **catdscfg** command to list the contents of the property files for the Device server and to check the status of the Device server. You must have Administrator authority to use this command.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is `system:port_number`, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing the contents of the property file

The following command lists the contents of the property file:

```
tpctool> catdscfg -url localhost:9550 -user ***** -pwd *****
```

The following output is returned:

Property	Context	Value
AgentManager.Registration	AM	YES
Fabric.Manager.Password	AM	*****
Fabric.Manager.Username	AM	manager
Registration.Server.Host	AM	tivoli12
Registration.Server.Port	AM	9511
Registration.Server.PW	AM	*****
TPCData.Password	AM	*****
TPCData.UserName	AM	manager
default.authorization.enabled	CIM	true
default.credential	CIM	default
default.principal	CIM	default
FabricCIMTransactionTimeout	CIM	600
FabricCIMZoneDBChecksumUsage	CIM	true
FabricCIMZSetActivationRetryInterval	CIM	20000
FabricCIMZSetActivationRetryMaxCount	CIM	9


Note: This is a partial sample of the actual output.

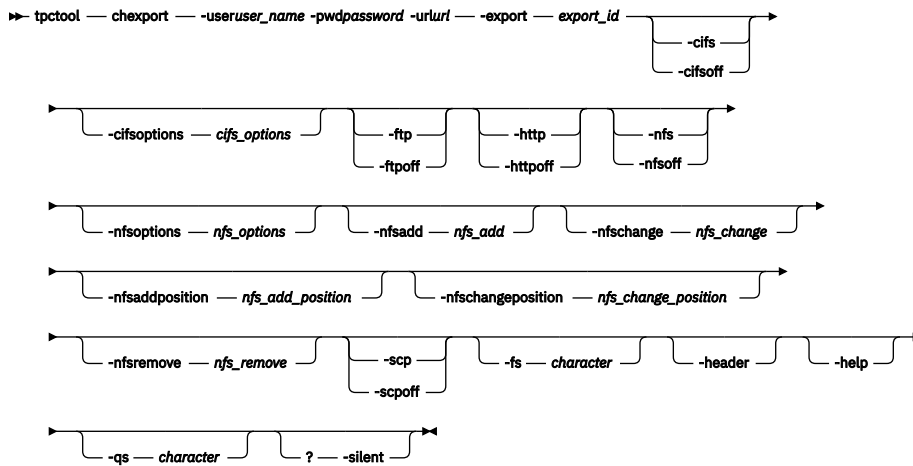
chexport

Use the **chexport** command to change the protocol configuration for an export.

An *export* is a shared disk space that is accessible through the protocols that you specify when you run the **mkexport** command. You can create exports and enable them for HTTP, FTP, Secure Copy Protocol (SCP), Network File System (NFS), and Common Internet File System (CIFS) protocols. You must have Administrator authority to use this command.

Syntax





Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- export export_id
Specifies the name of the export. This name is the export key that is listed in the ID column of the **lsexport** command output.
- active | -inactive
Marks the export as temporarily active or inactive. You can access the data only in an active export.
- cifs | -cifsoff
Adds or removes the CIFS protocol for the export.
- cifsoptions cifs_options
Defines the CIFS protocol options for the export. If the *cifsoptions* value contains spaces, the entire option must be enclosed in matching single quotation marks and the quotation marks must be preceded by an escape character.
- ftp | -ftppoff
Adds or removes FTP for the export.
- http | -httppoff
Adds or removes HTTP for the export.
- nfs | -nfsoff
Adds or removes the NFS protocol for the export.
- nfsoptions nfs_options
Defines the NFS clients and their options for the export.
- nfsadd nfs_add
Adds the NFS clients and their options to the export.
- nfschange nfs_change
Modifies the NFS clients and their options for the export.
- nfsaddposition nfs_add_position
Specifies the position of a new NFS entry that was specified by the *nfsadd* parameter.
- nfschangeposition nfs_change_position
Specifies the position of a modified NFS entry that was specified by the *nfschange* parameter.
- nfsremove nfs_remove
Removes one or more NFS clients from the export.
- scp | -scppoff
Adds or removes SCP for the export.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Changing the FTP and HTTP options for an export

The following command adds FTP and removes HTTP from the export.

```

tptool> chexport
-export cindyexport+storage1.storage.tucson.ibm.com+storage1.storage.tucson.
ibm.com+0 -user admin -pwd password -url localhost:9550 -ftp -httppoff

```

The following output is returned:

```
ExportId
=====
cindyexport+storage1.storage.tucson.ibm.com+storage1.storage.tucson.ibm.
com+0

Status
=====
SUCCESS
```

Example: Defining the CIFS options for an export

The following command defines the CIFS protocol options for the export.

```
chexport -export eexp10+kq98n5d.ibm+00000200A22045DC+0 -cifs
-cifsoptions "browseable=no,\"comment=comment for eexp10\",leases=no,
sharemodes=no,syncio=yes,hideunreadable=yes,cifsac1=no,oplocks=no,
locking=no,\"read only\",synconclose=no,\"access control=Everyone:ALLOWED:FULL;
Administrator:ALLOWED:FULL\"-user db2admin -pwd g0vmware -url localhost:9550
```

The following output is returned:

```
ExportId                      Status
=====
eexp10+kq98n5d.ibm+00000200A22045DC+0 SUCCESS
```

Example: Changing the NFS options for an export

The following command adds, modifies, and removes NFS entries for the export.

```
tpctool> unlinkfset -fileset eefset01+eefs+kq458mv.ibm+00000200A2A0153C+0
-user admin -pwd password -url localhost:9550
```

The following output is returned:

```
ExportId
=====
eexp600+kq98n5d.ibm+00000200A22045DC+0

Status
=====
SUCCESS
```

Example: Changing the NFS options and their positions for an export

The following command adds and modifies NFS entries and their positions for the export.

```
tpctool> chexport -export eexp700+kq98n5d.ibm+00000200A22045DC+0
-nfsadd "host4(ro)" -nfsaddposition host2 -nfschange "host2(ro)"
-nfschange position 3 -active -user db2admin -pwd g0vmware
-url localhost:9550
```

The following output is returned:

```
ExportId
=====
eexp700+kq98n5d.ibm+00000200A22045DC+0

Status
=====
SUCCESS
```

Related reference

- [mkexport](#)
- [lsexport](#)

Related information

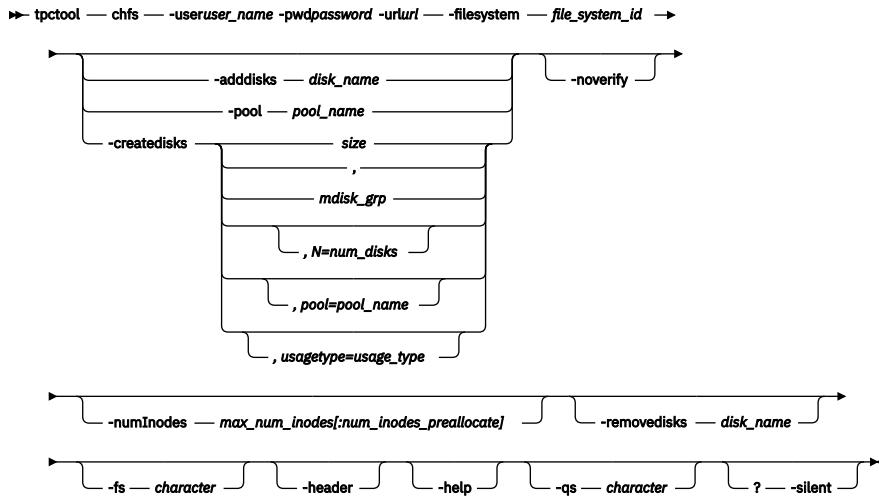
- <https://www.ibm.com/support/knowledgecenter/ST5Q4U>

chfs

Use the **chfs** command to change the properties of a file system, and to add or remove disks to or from a file system.

If you use this command to both add and remove disks, disks are added first and then disks are removed. Other settings are applied after disks are added or removed. You must have Administrator authority to use this command.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- filesystem file_system_id
Specifies the ID of file system to be changed.
- adddisks disk_name
Specifies the disks to add to the file system. The *disk_name* variable contains a comma-separated list of disk names.
Tip: You can verify the availability of a disk by running the **lsnsd** command.
- pool pool_name
Adds to the file system a set of free disks that have the file system pool name that is set as a storage pool. On IBM® Storwize® V7000 Unified, all disks in the file system pool must be tagged for this file system to be used.
Tip: You can list all of the disks in a storage pool by running the **lsnsd** command.
- createdisks size | mdisk_grp | N=num_disks | pool=pool_name | usage_type=usage_type
Creates disks implicitly, and then adds them to the file system. This option is applicable only for Storwize V7000 Unified.
 - size
Specifies the size of the new disks. Size is specified as an integer with capacity up to a petabyte without a space between the size and the unit; for example 17G. Disk sizes must be specified either without suffix (byte) or with K (kilobyte), M (megabyte), G (gigabyte), T (terabyte), or P (petabyte). Values less than 512 MB are not supported. This parameter is mandatory.
 - mdisk_grp
Specifies the storage system managed disk (MDisk) group in which the underlying NAS volumes are created. This parameter is mandatory.
Tip: You can see a list of available MDisk groups by using the **svcinfo lsmdiskgrp** command.
 - num_disks
Specifies the number of storage system NAS volumes that is created in each MDisk group. This parameter is optional. The default number of disks is 1.
 - pool
Specifies a pool for the disks. This parameter is optional. The default value is system.
 - usage_type
Specifies the usage type for the disks. This parameter is optional. Specify one of the valid usage types.
 - dataAndMetadata
 - dataOnly
 - metadataOnly

The default usage type is dataAndMetadata. The only valid usage type for a non-system pool is dataOnly.
- noverify
Specifies that disks must not be verified as belonging to an existing file system. If this option is used, either the -adddisks parameter or the -pool parameter must also be specified.
- numInodes max_num_inodes[:num_inodes_preallocate]
Specifies the maximum number of files for this file system. The *num_inodes_preallocate* variable specifies the number of inodes that the system immediately preallocates. You can specify values in thousands (k) or in millions (M). To specify values of 100 million for the *max_num_inodes* variable and 1 million for the *num_inodes_preallocate* variable, enter **-numInodes 100M:1M**.
- removedisks disk_name
Specifies the disks to remove from the file system. On IBM Storwize V7000 Unified, this option completely removes the specified disks and the data on them by deleting the volumes on the storage system.
Tip: You can verify the availability of a disk by running the **lsnsd** command.

- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Change a file system

The following command changes a file system by removing a disk.

```
tpctool> chfs -filesystem eefs+kq458mv.ibm+00000200A2A0153C+0
-user admin -pwd password -url localhost:9550 -removedisks IFS1319490696615
```

The following output is returned:

FilesystemId	Status
eefs+kq458mv.ibm+00000200A2A0153C+0	SUCCESS

Related reference

- [lsnsd](#)

Related information

- <https://www.ibm.com/support/knowledgecenter/ST5Q4U>

chfset

Use the **chfset** command to change the attributes of a fileset.

You can change the name of, or comment that is associated with, an existing file set. You can also change the maximum number of inodes and the number of inodes to allocate for an independent fileset. You must have Administrator authority to use this command.

Syntax

```
tpctool — chfset — -user user_name -pwd password -url url — -filesystem file_set_id — -name name —
- comment comment — -numInodes max_num_inodes[:num_inodes_preallocate] —
- fs character — -header — -help — -qs character — ? — -silent
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- filesystem file_set_id
Specifies the IBM Spectrum® Control key of the fileset to be changed. The fileset key is listed in the ID column of the **lsfset** command output.
- name name
Specifies the new name for this fileset.
- comment comment
Specifies a new comment that displays in the output of the **lsfset** command. The length of this comment can be a maximum of 255 characters. You must enclose comments in double quotation marks.
- numInodes max_num_inodes[:num_inodes_preallocate]
Specifies the maximum number of files for this fileset. The *num_inodes_preallocate* variable specifies the number of inodes that the system immediately preallocates. You can specify values in thousands (k) or in millions (M). To specify values of 100 million for the *max_num_inodes* variable and one million for the *num_inodes_preallocate* variable, enter *-numInodes 100M:1M*. GPFS defines a minimum number of inodes, which might be greater than the maximum specified. The default values for the fileset are one million (1M) for the *max_num_inodes* variable and 50,000 (50 K) for the *num_inodes_preallocate* variable.
- fs character

- Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Changing the attributes of a fileset

The following command changes the name of a fileset.

```
tpctool> chfset -fileset eefset+eefs+kq458mv.ibm+00000200A2A0153C+0
-name eefset01 -user admin -pwd password -url localhost:9550
```

The following output is returned:

FilesetId	Status
eefset01+eefs+kq458mv.ibm+00000200A2A0153C+0	SUCCESS

Related reference

- [lsfset](#)

Related information

- <https://www.ibm.com/support/knowledgecenter/ST5Q4U>

chkquota

Use the **chkquota** command to check quota limits for users, user groups, and filesets on a file system and to write the information to the GPFS database.

Because quota information in the database is not updated in real time, you must run the **chkquota** command before the **lsquota** command. This sequence refreshes the quota information in the GPFS database. You must have Administrator authority to use this command.

By specifying a file system, you can use the **chkquota** command to check the quotas for users, user groups, and filesets. This process can be lengthy because these devices can have large numbers of files. When the command completes, the collected data is updated in the GPFS database and is then available to the **lsquota** command for retrieval. The **chkquota** command output displays a completion status of success or error. It does not display the quota information. The **lsquota** displays the data that is collected by the **chkquota** command.

Syntax

```
tpctool — chkquota — -user user_name -pwd password -url url — -filesystem file_system_id →
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- filesystem file_system_id
Specifies the ID for the file system that is associated with the quota. You can use the **lsfs** command to return information, including the IDs, for all file systems that are discovered. The ID is listed in the ID column of the **lsfs** command output.

Example: Check a quota for a file system

The following command checks the status of a quota for a file system:

```
tpctool> chkquota -filesystem gpfs1+kq98n5d.ibm+00000200A16045DC+0
-user admin -pwd password -url localhost:9550
```

The following output is returned:

FilesystemId	Status
gpfs1+kq98n5d.ibm+00000200A16045DC+0	SUCCESS

Related reference

- [lsquota](#)
- [lsfs](#)

Related information

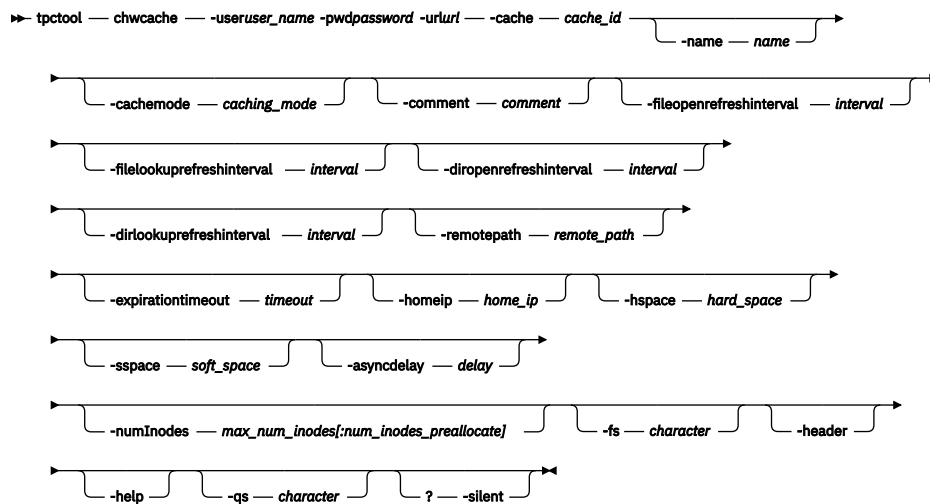
- <https://www.ibm.com/support/knowledgecenter/ST5Q4U>

chwcache

Use the **chwcache** command to modify the attributes of a cache fileset on a wide area network (WAN) cluster.

Data on a home system is cached to this fileset. Use the **lswcache** command to view the attributes of the cache fileset. You must have Administrator authority to use this command.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- cache cache_id
Specifies the ID of the cache fileset to be changed.
You can use the **lswcache** command to retrieve the cache fileset ID.
- name name
Specifies the new name for the cache fileset.
- cachemode read-only | local-updates | single-writer
Specifies the operating mode for the cache system. The default mode is *read-only*.
- comment comment
Specifies a comment for the cache.
- fileopenrefreshinterval interval
The maximum interval, in seconds, between when a file is opened on the cache system and when it was last validated with the home system. The default value is 30.
- filelookuprefreshinterval interval
The maximum interval, in seconds, between when a file is accessed on the cache system and when it was last validated with the home system. The default value is 30.
- directoryopenrefreshinterval interval
The maximum interval, in seconds, between when a directory is opened on the cache system and when it was last validated with the home system. The default value is 60.
- directorylookuprefreshinterval interval
The maximum interval, in seconds, between when a directory is accessed on the cache system and when it was last validated with the home system. The default value is 60.
- remotepath remote_path

Alternatively, you can specify the path of the cache source, and the IP address is determined automatically. For example, `-remotepath /ibm/gpfs0/wcacheSource`.

-expirationtimeout timeout

-homeip home_ip

Restriction: You cannot use FQDNs.

Specifies the hard limit or maximum of disk space usage by the fileset created for the WAN cache. The default value is 0, which implies there is no limit.

Specifies the soft limit or minimum of disk space usage by the fileset created for the WAN cache. The default value is 0, which implies there is no limit.

The time interval, in seconds, between a write operation on the cache filesset and the corresponding update on the home filesset. The default value is 15.

Defines the inode limits for the fileset created for the WAN cache. The *max_num_inodes* variable specifies the maximum number of inodes that can be allocated to the fileset. The *num_inodes_preallocate* variable specifies the number of inodes that the system immediately preallocates. You can specify values in thousands (K) or in millions (M). To specify values of 100 million for the *max_num_inodes* variable and 1 million for the *num_inodes_preallocate* variable, enter `-numInodes 100M:1M`. If not provided the default values are 100K:100K.

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

Lists help information for the command.

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").

Suppresses all output for the command. If you omit this parameter, output is included.

The following command changes the attributes of a cache fileset.

The following output is returned:

CacheId	Status
ctest+ee+storage1.storage.tucson.ibm.com+127.0.0.1+0	SUCCESS

Use the **chwcachesource** command to change the properties of a home system on a fileset. The home system is the source of the data in a wide area network (WAN) cache configuration. You can use WAN caching to distribute data transparently among data centers and multiple remote locations without disruption to applications. You must have Administrator authority to use this command. You can modify the properties of the home system by adding or removing cache systems. You can also update the public keys on all the cache systems for the specified home system.

```
➤ tpctool — chwcachesource — -useruser_name -pwdpassword -urlurl — -cachesource — cache_source_id ➔
```

┌───────────┴───────────┐ ┌───────────┴───────────┐ ┌───────────┴───────────┐

```
      -addclient — client_definition      -removeclient — client_cluster_id      -updatekeys
```

┌───────────┴───────────┐ ┌───────────┴───────────┐ ┌───────────┴───────────┐

```
      -fs — character      -header      -help      -as — character      ? — -silent
```

-user user_name	Specifies an IBM Spectrum Control user ID.
-pwd password	Specifies the password for the IBM Spectrum Control user ID.
-url url	Specifies the Device server. The format of the URL is <i>system:port_number</i> , where <i>system</i> represents either the host name or IP address, and <i>port_number</i> represents the IBM Spectrum Control Device server port.
-cachesource cache_source_id	

- Specifies the ID of the home system that is created by the **mkwccachesource** command. Use the **lswccachesource** command to retrieve the cache source ID.
- addclient *client_definition*
Adds a cache system to the existing list of clients for a specified home system. The *client_definition* variable contains a comma-separated list of the IP address and access mode of the management nodes for the cache systems. Access mode can be either ro (read-only) or rw (read/write). Only one of the cache clients can have read/write permission at any time. You must enclose this parameter in single quotation marks.
Restriction: You cannot add a cache client with read/write access if there is an export with read/write access for the same path. For example, an export is enabled on protocol types like NFS, CIFS, or FTP.
- removeclient *client_cluster_id*
Removes a cache system from the existing list of clients for a specified home system. The *client_cluster_id* variable is the cluster ID of the cache system. Use the **lswccachesource** command with the **-l** option to find the cluster ID of a specified cache system.
Restriction: You cannot remove the last cache system for a specified cache home system.
- updatekeys
Updates the public keys of all the cache systems for the specified home systems.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Adding a cache system to a home system that is enabled for WAN caching

The following command adds a cache system to a home system on the home cluster.

```
tpctool> chwccachesource -cachesource eesrc10+storage1.storage.tucson.
ibm.com+127.0.0.1+0 -addclient '127.0.0.2(ro)'
```

The following output is returned:

CachesourceId	Status
===== eesrc10+storage1.storage.tucson.ibm.com+127.0.0.1+0	SUCCESS

Example: Removing a cache system from a home system that is enabled for WAN caching

The following command removes a cache system from a home system on the home cluster.

```
tpctool> chwccachesource -cachesource eesrc10+storage1.storage.tucson.
ibm.com+127.0.0.1+0 -removeclient 792217928950257960
```

The following output is returned:

CachesourceId	Status
===== eesrc10+storage1.storage.tucson.ibm.com+127.0.0.1+0	SUCCESS

ckzone

Use the **ckzone** command to verify that a fabric contains a zone.

Syntax

```
➔ tpctool — ckzone — -user user_name -pwd password -url url — -fabric WWN — zone —>
                                     -help ? -silent
```

Parameters and arguments

- user *user_name*
Specifies an IBM Spectrum Control user ID.
- pwd *password*
Specifies the password for the IBM Spectrum Control user ID.
- url *url*
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric *WWN*
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- help | -h | -?

- Lists help information for the command.
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- zone
Specifies the zone.

Example: Verifying that a fabric contains a zone

The following command checks whether the fabric contains the *SUNSOLARIS* zone.

```
tpctool> ckzone -user me -pwd mypass -url myhost:myport
-fabric 100000051E34F6A8 SUNSOLARIS
```

If the fabric contains the zone, the following text is returned:

```
Zone SUNSOLARIS found in fabric 100000051E34F6A8
```

ckzs

Use the **ckzs** command to verify that a fabric contains a zone set.

Syntax

```
tpctool — ckzs — -user user_name -pwd password -url url — -fabric — WWN — zone_set
```

-help ? — -silent

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- help | -h | -?
Lists help information for the command.
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- zone_set
Specifies the zone set. The *zone_set* variable is the name of the zone set.

Example: Verifying that a fabric contains a zone set

The following command determines whether the fabric contains the *PARIS* zone set.

```
tpctool> ckzs -user me -pwd mypass -url myhost:myport -fabric 100000051E34F6A8 PARIS
```

If the fabric contains the zone set, the following text is returned:

```
Zoneset PARIS found in fabric 100000051E34F6A8
```

commit

Use the **commit** command to commit a transaction. When you commit a transaction, all the commands issued after you started the transaction are enacted. You must have Administrator authority to use this command.

Syntax

```
tpctool — commit — -user user_name -pwd password -url url — -fabric — WWN —
```

-help

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fabric WWN

Specifies the fabric. The *WWN* variable is the worldwide name (WWN).

-help | -h | -?

Lists help information for the command.

Example: Committing a transaction

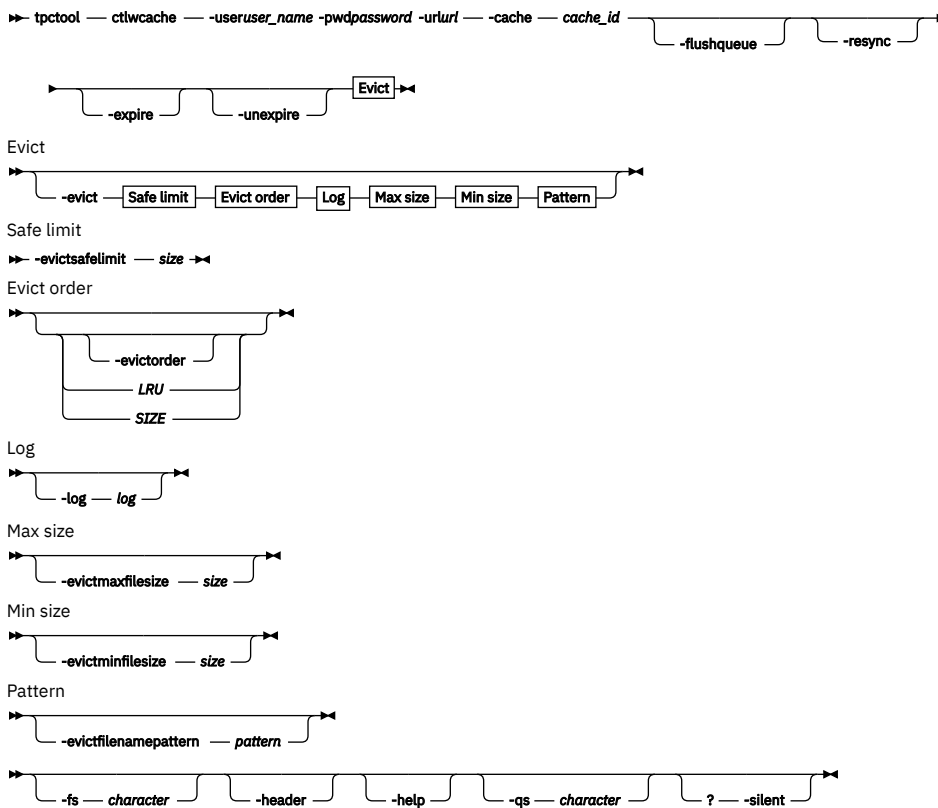
The following command commits a transaction. The user has previously specified connection options, started a transaction, and issued a fabric-control command:

```
tpctool> commit -fabric 100000051E34F6A8
```

ctlwcache

Use the **ctlwcache** command to run maintenance operations on a cache fileset on a wide area network (WAN) cluster. Data on a source cluster is cached to this fileset. You can resynchronize cache data to the source fileset, flush the cache to the source fileset, remove cached data, and expire or unexpire the data in a fileset. Use the **lswcache** command to view the attributes of the cache fileset. You must have Administrator authority to use this command.

Syntax



Parameters and arguments

-user user_name

Specifies an IBM Spectrum Control user ID.

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-cache cache_id

Specifies the ID of the cache fileset to be changed. You can use the **lswcache** command to retrieve the cache fileset ID.

-name name

Specifies the new name for the cache fileset.

-flushqueue

Flushes the queue at the cache gateway node to the source file set. The command completes when the queue is empty.

-resync

Synchronizes the cache source and the contents of the cache file set. Various events can cause inconsistencies between the source and cache versions of data. Use this parameter to ensure that the source version is consistent with the cache version.

Restriction: You can use this option only when the cache state is active.

- expire
Marks all the cache fileset contents as expired.
You must ensure that the `expirationtimeout` setting for the cache fileset is enabled. To view the `expirationtimeout` value, use the **lswcache** command with the `-v` parameter. You can change the `expirationtimeout` value by using the **chwcach** command.
- unexpire
Removes the expired status from all of the cache fileset contents that are marked as expired.
You must ensure that the `expirationtimeout` setting for the cache fileset is enabled. To view the `expirationtimeout` value, use the **lswcache** command with the `-v` parameter. You can also change the `expirationtimeout` value by using the **chwcach** command.
- evict
Removes cache data if the soft limit is set for the disk usage by the fileset. You can remove cache data either when the cache file set is created by using the **mkwcache** command, or later by using the **setquota** command. You can use the **lsquota** command to view quota values. If you specify the `-evict` parameter, you must specify the `-evictsafelimit` parameter.
- evictsafelimit size
Specifies the safe limit for eviction. This safe limit overrides the soft limit value during the manual eviction operation. You must specify a value that is less than the value for the soft limit. Specify the size as an integer either without suffix (byte) or with `K` (kilobyte), `M` (megabyte), `G` (gigabyte), `T` (terabyte), or `P` (petabyte). These suffixes are not case-sensitive.
- evictorder LRU | SIZE
Evicts the contents of a specified cache fileset based on the order of the queue. Specify `LRU` to evict least recently used files first. Specify `SIZE` to evict larger-sized files first.
- log log
Specifies the location of the eviction log file. The log file location must be within an active export. By default, the logged information is appended to the system event log, `mmfs.log`.
- evictmaxfilesize size
Specifies the maximum file size for eviction. Specify the size as an integer (byte) or with the suffix `K` (kilobyte), `M` (megabyte), `G` (gigabyte), `T` (terabyte), or `P` (petabyte). These suffixes are not case-sensitive.
- evictminfilesize size
Specifies the minimum file size for eviction. Specify the size as an integer (byte) or with the suffix `K` (kilobyte), `M` (megabyte), `G` (gigabyte), `T` (terabyte), or `P` (petabyte). These suffixes are not case-sensitive.
- evictfilenamepattern pattern
Specifies the file name pattern for eviction. You can use the percent sign (%) as the wildcard character.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Synchronizing a cache source and the contents of the cache fileset

The following command synchronizes a cache source and the contents of the cache fileset.

```
tpctool> ctlwcache -cache cache3+ad+storage3.storage.tucson.
ibm.com+127.0.0.1+0 -resync
```

The following output is returned:

CacheId	Status
cache3+ad+storage3.storage.tucson.ibm.com+127.0.0.1+0	SUCCESS

Example: Removing a cache fileset

The following command removes the contents of a specified cache fileset based on size order, by using an eviction safe limit of 10,000 bytes.

```
tpctool> ctlwcache -cache cache3+ad+storage3.storage.tucson.
ibm.com+127.0.0.1+0 -evict -evictsafelimit 10000 -evictorder SIZE
```

The following output is returned:

CacheId	Status
cache3+ad+storage3.storage.tucson.ibm.com+127.0.0.1+0	SUCCESS

Related reference

- [chwcach](#)
- [lsquota](#)
- [setquota](#)

deactzs

Use the **deactzs** command to deactivate the active zone set. This command must be run as a transaction. You must have Administrator authority to use this command.

Syntax

➤ **tpctool** — **deactzs** — **-user***user_name* **-pwd***password* **-url***url* — **-fabric** — *WWN* — **-help** — **? — -silent** ➤

Parameters and arguments

- user** *user_name*
Specifies an IBM Spectrum Control user ID.
- pwd** *password*
Specifies the password for the IBM Spectrum Control user ID.
- url** *url*
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric** *WWN*
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- help** | **-h** | **-?**
Lists help information for the command.
- silent**
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Deactivating the active zone set

The following commands deactivate the active zone set:

```
tpctool> start -user me -pwd mypass -url myhost:myport -fabric 100000051E34F6A8
tpctool> deactzs -fabric 100000051E34F6A8
tpctool> commit -fabric 100000051E34F6A8
```

encrypt

Use the **encrypt** command to generate an encrypted password for use in the configuration file. This command takes text from standard input and generates 7-bit ASCII-equivalent characters (uencode).

Syntax

➤ **tpctool** — **encrypt** — *password* ➤

- password*
Specifies the password to be encrypted.

Example: Encrypting a password

The following command encrypts the specified password:

```
tpctool encrypt myverylongpassword
```

getdscfg

Use the **getdscfg** command to list the current value of a property from the property file for the Device server. You must have Administrator authority to use this command.

Syntax

➤ **tpctool** — **getdscfg** — **-user***user_name* **-pwd***password* **-url***url* — **-property** — *property_key* — **-context** — *context* — **-fs** — *character* — **-header** — **-help** — **-qs** — *character* — **? — -silent** ➤

- user** *user_name*
Specifies an IBM Spectrum Control user ID.
- pwd** *password*

- url url
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- property property_key
Specifies the property key. The *property_key* variable is the property key.
- context context
Specifies a classification or category for a configuration property. The *context* variable is the context properties. For example, The following parameter applies to the IBM Spectrum® Control Device server only.
-context DeviceServer
- The following parameter applies to the IBM Spectrum Control performance manager only.
-context PerformanceManager
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing a property value

The following command lists the value of the SnmpRetryCount property:

```
tpctool> getdscfg -url localhost:9550 -user ***** -pwd ***** -property
SnmpRetryCount -context DeviceServer
```

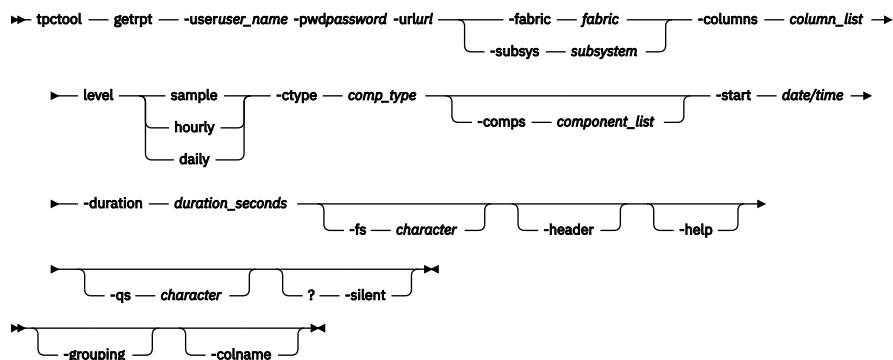
The following output is returned:

Property	Context	Value
SnmpRetryCount	DeviceServer	3

getrpt

Use the **getrpt** command to list a performance report for a specified storage subsystem.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric fabric
Specifies the name of a switch, such as one that is returned by the **lsdev** command.
- subsys subsystem
Specifies the name of a storage subsystem, such as one that is returned by **lsdev**. The name of the storage subsystem is the globally unique identifier (GUID) of that storage subsystem.

- columns *column_list*
Specifies what columns displays in the report. The columns are obtained from the **lscounters** and **lsmetrics** commands.
- level *sample | hourly | daily*
Specifies the level for which the performance metrics that are summarized. You can specify a sample summary, an hourly summary, or a daily summary.
- ctype *comp_type*
Specifies that the output includes only components of the specified type. For more information about the *comp_type* variable, see the **lstype** command.
- comps *component_list*
Specifies the component list. The *component_list* variable specifies the components, such as one that is returned by **lscomp**.
- start *date/time*
Specifies the date and time to start the sampling period. The date and time are formatted as:

yyyy.MM dd:HH:mm:ss

All time zones are relative to the Device server. For more information, see the **lstime** command.
- duration *duration_seconds*
Specifies the duration of the sampling period, in seconds. The *duration_seconds* variable is an integer.
- fs *character*
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs *character*
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- grouping
Enables the grouping of numeric values. For example, in English the value 12000 would display as 12,000. The grouping character is determined by the system locale.
- colname
Displays the name of the counter or metric in the column header. By default, an integer that indicates the column or metric type is displayed in the column header.

Example: Listing performance metrics

The following command lists a report of performance metrics:

```
tpctool> getrpt -user me -pwd mypass -url myhost:myport
-fabric 100005668 -subsys 2105.22232+0 -level daily
```

Note: If you have multiple volumes in your subsystems, use the -fs *character* option with a comma as a field separator.

```
tpctool> getrpt -subsys 2107.1302541+0 -columns 1,2,4,5,10,11,13,14,22,23,28,29,31,
32,33,34,35,36,37,38,39,40,41,42,43,44,45,46,47,48,49,50,51,52,53,54,55,56,57,58,5,
9,828,829,830,831,832,833,801,802,803,804,805,806,807,808,809,810,811,812,813,814,
815,816,817,818,819,820,821,822,823 -ctype vol -level sample -start 2007.06.27:19:
01:40 -duration 3600 -fs
```

Related reference

- [lstype](#)

linkset

Use the **linkset** command to create a junction to connect a name in a directory of a parent file set to the root directory of a child fileset. You must have Administrator authority to use this command.

As a prerequisite, the file system must be mounted and the junction path must be under the mount point of the file system. A dependent file set that is contained within an owner fileset can be linked only inside its owner path or in a directory tree of the owner path. A fileset without an owner can be linked anywhere in the file system.

Although a junction is displayed as a directory, you cannot issue directory commands to remove a junction. Use the **unlinkset** command instead.

Syntax

```
➔ tpctool — linkset — -user user_name -pwd password -url url — -help — -fileset — file_set_id — -path — path —
-headers — ? — -silent — -qs — character — -fs — character —
```

Parameters and arguments

- user *user_name*
Specifies an IBM Spectrum Control user ID.
- pwd *password*
Specifies the password for the IBM Spectrum Control user ID.
- url *url*

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-help | -h | -?

Lists help information for the command.

-fileset file_set_id

Specifies the IBM Spectrum® Control key of the fileset to be linked. The fileset key is listed in the ID column of the **lsfset** command output.

-path path

Specifies the name of the junction. The name must not refer to an existing file system object.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

Example: Linking a fileset

The following command links the fileset that is named **eefset01+eefs+kq458mv.ibm+00000200A2A0153C+0** to the path **/ibm/eefs/eefsetdir01**.

```
tpctool> linkset -fileset eefset01+eefs+kq458mv.ibm+00000200A2A0153C+0  
-path /ibm/eefs/eefsetdir01 -user admin -pwd password -url localhost:9550
```

The following output is returned:

FilesetId	Status
eefset01+eefs+kq458mv.ibm+00000200A2A0153C+0	SUCCESS

Related reference

- [unlinkset](#)
- [lsfset](#)

Related information

- <https://www.ibm.com/support/knowledgecenter/ST5Q4U>

lsappgroup

Use the **lsappgroup** command to display a list of applications that are known to IBM Spectrum Control.

Syntax

```
tpctool — lsappgroup — -user user_name -pwd password -url url  
                                -fs character -header -help  
-qs character -? -silent -l
```

Parameters and arguments

-user user_name

Specifies an IBM Spectrum Control user ID.

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").

-l

Specifies the long version of the information:

- Name
Specifies the name that identifies the application.
- ID
The unique ID of the application. For example, **/Application/5044** where 5044 is the database ID of the application.
- Parent(s)
Specifies the direct parent of the application. For example, **Db2**.
- Type
Specifies the type of application.
- Subtype
Specifies another qualifier for the type of application. For example, if **HRApplication** is a type of application then **Project** might be a subtype.
- Status
Specifies the health of the application. For example, **NORMAL**, **WARNING** or **ERROR**.
- Description
Specifies the description of the application.
- User-defined property 1, 2, 3
A user-defined property (UDP1,2,3) that can contain anything of significance to the administrator.
- Path(s)
Specifies the hierarchy path of the application from the top level on down. For example, **/Lab/Testlab**.

Example: Displaying a list of applications

The following command generates a list of applications that are known to IBM Spectrum Control:

```
tpctool>lsappgroup
```

The following output is returned:

Name	ID	Parents	Type
Appgroup1	/Application/5035		HRApplication
Appgroup2	/Application/5041		Medical
Db2appgroup	/Application/5044	Db2	
Accounting	/Application/5046		Finance
Testlab	/Application/5052	"Lab, Test"	TestLab

The following command generates a list of applications that are known to IBM Spectrum Control, and the long version of the information:

```
tpctool>lsappgroup -l
```

The following output is returned:

Name	ID	Parents	Type	Subtype	Status	Description
UDP1	UDP2	UDP3	Paths			
Appgroup1	/Application/5035		HRApplication	Project	ERROR	"Description 1"
Appgroup2	/Application/5041		Medical	Program	ERROR	"Description 2"
Db2appgroup	/Application/5044	Db2	IT	Project	ERROR	"Description 3"
Accounting	/Application/5046		Finance	Program	NORMAL	"Description 4"
Testlab	/Application/5052	"Lab, Test"	TestLab	-	ERROR	-
UDP1	UDP2	UDP3	Paths			
-	-	-	/Appgroup1			
attrib1	-	-	/Appgroup2			
attrib1-	-	-	/Db2/Db2appgroup			
-	-	-	/Accounting			
-	-	-	"/Lab/Testlab, /Test/Testlab"			

Related reference

- [lsappgroupmembers](#)

lsappgroupmembers

Use the **lsappgroupmembers** command to list members of a specified application.

Syntax

```
tpctool -- lsappgroupmembers -- -user user_name -pwd password -url url
                                     -fs character -header
                                     -id ApplicationGroupID
                                     -help -qs character ? -silent -l
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- id ApplicationGroupID
The unique ID of the application. For example, `/Application/5044` where **5044** is the database ID of the application. The output includes the following information:
 - Member Key
 - Member Name
 - Member Type
- l
Specifies that the long version of the information is listed. The following information is included:
 - UDP1
The user-defined property (UDP) contains anything of significance to the group administrator. This alphanumeric character can contain any alphanumeric character, and must be fewer than 255 characters in length.
 - UDP2
The user-defined property (UDP) contains anything of significance to the group administrator. This alphanumeric character can contain any alphanumeric character, and must be fewer than 255 characters in length.
 - UDP3
The user-defined property (UDP) contains anything of significance to the group administrator. This alphanumeric character can contain any alphanumeric character, and must be fewer than 255 characters in length.

Example: Displaying a list of the members for an application

The following command generates a list of the members for a specified application:

```
tpctool> lsappgroupmembers -id /Application/5044 -l
```

The following output is returned.

Member Key	Member Name	Member Type
george_ifs_share2+kq98n5d.ibm+00021	george_ifs_share2	Export
omni+Atlas.storage.abc.ibm.com	omni	Virtual Machine
XIV_1+Atlas.storage.abc.ibm.com	XIV_1	Data Store
root+SMAC+kq99n5d.ibm+00000200A3C044DC+0	root	Fileset

UDP1	UDP2	UDP3
ABCLab	-	-
ABCLab	-	-
-	-	-
-	-	-


Related reference

- [lsappgroup](#)

Isarray

Use the **lsarray** command to display information about arrays and back-end storage systems. For SAN Volume Controller, XIV® systems, and Dell EMC VMAX and VNX storage systems, the command displays information about block pools. You use the **setarray** command or **autosetarray** command to provide information about back-end storage systems. You can also enter information about back-end storage systems on the Pool Properties page in the IBM Spectrum Control GUI.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- dev system
Specifies the globally unique identifier (GUID) of the storage system. You can use the **lsdev** command to return information, including the GUID, for all storage systems that are discovered by IBM Spectrum Control.
- l
Specifies that the detailed version of information about arrays is provided. If you issue the **lsarray** command and do not specify the -l parameter, only the array ID is displayed.

Column label	Details
Array	The ID of the array.
Label	The name of the storage pool (managed disk group).
Total Size (GB)	The total amount of storage space.
Free Size (GB)	The total amount of available storage space.
Status	The operational status of the storage pool that hosts the array.
Types of Disk	The types of disk that host the arrays. <ul style="list-style-type: none"> • <code>Solid State</code> for solid-state disks • <code>Non-Solid State</code> for hard disks • <code>Mixed</code> for solid-state disks and hard disks
Encrypted	Shows whether arrays are hosted on encrypted disks.
Read I/O Capability	Shows the read I/O capability in seconds of storage arrays. The calculation of read I/O capability is based on the following characteristics. <ul style="list-style-type: none"> • The type of storage system • The type of RAID • The type of disk • The number of disks
Back-end Type	The type of back-end storage system that manages disk group storage.
Back-end RAID Type	The type of Redundant Array of Independent Disk (RAID) that the back-end storage system uses.
Back-end Disk Type	The type of disk that the back-end storage system uses.
Back-end Disk Count	The total number of disks.

Tip: To find out how to set values for back-end storage subsystems, see the description of the **setarray** command.

- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- grouping
Enables the grouping of numeric values. For example, in English the value 12000 would display as 12,000. The grouping character is determined by the system locale.

Example: Listing array information

The following command lists IDs for the arrays on the specified system:

```
tpctool> lsarray -user me -pwd mypass -url myhost:9550 -dev 2107.75DG000+0
```

The following output is returned:

```
Array
=====
IBMTSDS: IBM.2107-75DG000-P17+2107.75DG000+0
IBMTSDS: IBM.2107-75DG000-P16+2107.75DG000+0
IBMTSDS: IBM.2107-75DG000-P15+2107.75DG000+0
IBMTSDS: IBM.2107-75DG000-P14+2107.75DG000+0
IBMTSDS: IBM.2107-75DG000-P1+2107.75DG000+0
IBMTSDS: IBM.2107-75DG000-P0+2107.75DG000+0
```

Example: Listing detailed information about an array

The following command provides detailed information about the array on the specified system:

```
tpctool> lsarray -dev 0000020064405BA0+0 -l
```

The following output is returned:

```
Array                               Label      Total Size (GB) Free Size (GB)
=====
0000020064405BA0:0+0000020064405BA0+0 mdisk_grp  407.38      310.25

Status  Types of Disk  Encrypted Read I/O Capability Back-end Type
=====
NORMAL  Non-Solid State  No      0      D

Back-end RAID Type Back-end Disk Type Back-end Disk Count
=====
5                  F10                  160
```

lsbackenddisktypes

Use the **lsbackenddisktypes** command to list the types of back-end disks and their average input/output.

This command is available for the following storage systems:

- Storwize® V7000
- SAN Volume Controller

Syntax

```
tpctool — lsbackenddisktypes — -user user_name -pwd password -url url
                                     -fs character -header
                                     -help -qs character ? -silent -grouping
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- grouping
Enables the grouping of numeric values. For example, in English the value 12000 would display as 12,000. The grouping character is determined by the system locale.

Example: Listing back-end disk types

The following command lists the types of back-end disks:

```
tpctool> lsbackenddisktypes
```

If the command is successful, a list of the types of back-end disks is displayed.

```
Back-end Disk Type  Description                               IOPS
=====
A07                  "Sata - 7 500 rpm"                        40
F10                  "Fiber - 10 000 rpm"                      120
F15                  "Fiber - 15 000 rpm"                      150
DEFAULT              "DEFAULT TYPE FOR UNCONFIGURED DEVICES"  0
```

When you issue the **lsbackenddisktypes** command, the following information is displayed.

Backend Disk Type
The type of back-end disk. For example: A07

Description

A description of the type of back-end disk. For example: "Sata - 7500 rpm"

IOPS

Input/Output (I/O) per second. The average number of input/output operations per second for the disk.

Tip: Use the **setbackenddisktype** command to set or update the type of a back-end disk.

lsbackendraidtypes

Use the **lsbackendraidtypes** command to list the types of back-end RAID arrays that are available for managed disk groups.

This command is available for the following storage systems:

- Storwize® V7000
- SAN Volume Controller

Syntax

```
tpctool lsbackendraidtypes -user user_name -pwd password -url url -fs character -header -help -qs character -silent -grouping
```

Parameters and arguments

- user** user_name
Specifies an IBM Spectrum Control user ID.
- pwd** password
Specifies the password for the IBM Spectrum Control user ID.
- url** url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fs** character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header**
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help** | **-h** | **-?**
Lists help information for the command.
- qs** character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent**
Suppresses all output for the command. If you omit this parameter, output is included.
- grouping**
Enables the grouping of numeric values. For example, in English the value 12000 would display as 12,000. The grouping character is determined by the system locale.

Example: Listing back-end RAID types

The following command lists the available types of back-end RAID types.

```
tpctool> lsbackendraidtypes
```

If the command is successful, a list of the available types of back-end RAID types is displayed.

```
Back-end RAID Type Weighted IO
=====
X                  1
6                  6
5                  4
1                  2
DEFAULT           1
```

When you issue the **lsbackendraidtypes** command, the following information is displayed.

Backend RAID Type

The type of back-end RAID array. Possible values are DEFAULT, 1, 5, 6, X, or an array type that is set with the **setbackendraidtype** command.

Weighted IO

The weighted input/output of the RAID type.

Tip: Use the **setbackendraidtype** command to set or update the type of a back-end RAID array.


lsbackendtypes

Use the **lsbackendtypes** command to list the types of back-end storage systems.

More information, such as the name, description, and cache hit ratio of the back-end type of the storage system is also provided. This command is available for the following storage systems:

- Storwize® V7000
- Storwize V7000 Unified
- SAN Volume Controller

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- grouping
Enables the grouping of numeric values. For example, in English the value 12000 would display as 12,000. The grouping character is determined by the system locale.

Example: Displaying a list of types of back-end storage systems

The following command provides a list of the types of back-end storage systems:

```
tpctool> lsbackendtypes
```

A list of the types of back-end storage subsystem is displayed.

Back-end Type	Name	Cache Hit Ratio
Fake	F800	100
X	XIV	70
D	DS8000	50
S	"EMC Symmetrix"	50
C	"EMC Clariion"	50
DEFAULT	DEFAULT	0

```
Description
=====
"Theoretical All Cache subsystem"
"IBM XIV"
"DS8K Disk Controller"
"DMX BE Cache Hit Ratio"
"CX BE Cache Hit Ratio"
"DEFAULT TYPE FOR UNCONFIGURED DEVICES"
```

When you issue the **lsbackendtypes** command, the following information is displayed:


- Back-end Type
The type of back-end storage system such as *D*
- Name
The name that is associated with the type of back-end storage system. For example, the name *DS8000* is associated with type *D*
- Cache Hit Ratio
The estimated cache hit ratio for read operations
- Description
A description of the type of back-end storage system

Tip: Use the **setbackendtype** command to set or update the type of back-end storage system.

Iscluster

Use the **lscluster** command to list all clusters or specified clusters that are on a Storwize® V7000 Unified storage system.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- dev dev_id | -cluster cluster_id...
Specifies that information is displayed for all clusters that are on a storage system. This parameter requires the globally unique identifier (GUID) of the storage system. You can use the **lsdev** command to return information, including the GUID, for all storage systems that were discovered.
The -cluster parameter specifies that information is displayed only for specific clusters. You can enter one or multiple cluster IDs. Use a comma to separate multiple IDs.
- l
Specifies that the following information is listed.

ID
An ID for the cluster that is composed of the cluster name, storage system name, and storage system format.
Cluster
The name of the cluster.
Cluster ID
The cluster ID as defined by the Storwize V7000 Unified file module.
Type
One or more of the following types: interface, management, or storage.

If you omit this parameter, only the ID is returned.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing all clusters that are on a storage system

The following command generates a list of all clusters that are on a Storwize V7000 Unified storage system.

```
tpctool> lscluster -dev 00000200A0E0005C+0 -l
```

The following output is returned:


ID	Cluster
storage2.storage.tucson.ibm.com+00000200A0E0005C+0	storage2.storage.tucson.ibm.com

Cluster ID	Type
12402779238946656959	interface,storage

lscomp

Use the **lscomp** command to list the components for which performance data is collected.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- subsys subsystem
Specifies the storage subsystem. The subsystem variable is the GUID of the storage subsystem. You can use the **lsdev** command to return information, including the GUID, for all storage subsystems that are discovered by IBM Spectrum® Control.
- ctype comp_type
Specifies that the output is to include only components of the specified type. For more information about the *comp_type* variable, see the **lstype** command.
- level sample | hourly | daily
Specifies the level for which the performance counters are to be summarized. You can specify a sample summary, an hourly summary, or a daily summary.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing performance counters

The following command generates a list of performance counters for volumes on the specified storage subsystem:

```
tpctool> lscounters -user me -pwd mypass -url myhost:myport
-subsys 2105.22232+0 -ctype vol -level sample
```

The following output is returned.

Counter	Value
Read I/O Count (normal)	1
Read I/O Count (sequential)	2
Write I/O Count (normal)	4
Write I/O Count (sequential)	5
Read Cache Hit Count (normal)	10
Read Cache Hit Count (sequential)	11
Write Cache Hit Count (normal)	13
Write Cache Hit Count (sequential)	14
Read Data Count	22
Write Data Count	23
Read Service Periods	28
Write Service Periods	29
Disk to Cache Transfers (normal)	31
Disk to Cache Transfers (sequential)	32
Cache to Disk Transfers	33
NVS Allocation Count	34
DFW I/O Count (normal)	35
DFW I/O Count (sequential)	36
NVS Delayed I/O Count	37
Cache Delayed I/O Count	38
Record-Mode-Read I/O Count	39
Record-Mode-Read Cache Hit Count	40
Quick Write Promote Count	41

Related reference

- [lstype](#)

Syntax



If the `-l` parameter is not issued, only the **Name**, **ID**, **Parents** and **Type** columns are returned.

Example: Displaying a list of departments

The following output is returned:

Name	ID	Parents	Type
Sales	/Department/5035	-	-
Accounting	/Department/5041	Finance	-
Tax	/Department/5044	Accounting,Sales	-
HR_WorldWide	/Department/5045	-	HR
HR_US	/Department/6245	HR_WorldWide	HR
HR Asia	/Department/6246	HR_WorldWide	HR

The following command generates a list of departments that are known to IBM Spectrum Control, and the long version of the information:

```
tpctool>lsdeptgroup -l
```

The following output is returned:

Name	ID	Parent(s)	Type	Subtype	Status
Sales	/Department/5035	-	-	Project	ERROR

Accounting	/Department/5041	-	Finance	Project	ERROR
Tax	/Department/5044	Accounting, Sales	-	Project	ERROR
HR_US	/Department/6245	HR_WorldWide	HR	Program	NORMAL
HR_Asia	/Department/6246	HR_WorldWide	HR	Program	Normal

Description	UDP1	UDP2	UDP3	PATHS
"Description 1"	-	-	-	/Department/5035
"Description 2"	-	-	-	-
"Description 3"	-	-	-	-
"Description 4"	-	-	-	-
"Description 5"	-	-	-	/Department/6246

Related reference

- [lsdeptgroupmembers](#)

lsdeptgroupmembers

Use the **lsdeptgroupmembers** command to list members of a specified department.

Syntax

```
tpctool — lsdeptgroupmembers — -user user_name -pwd password -url url
                                     -fs character -header
                                     -help -qs character ? -silent -id DepartmentGroupName -l
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- id DepartmentGroupName
The unique name of the department. In addition to the **Member ID** column, the output includes the following information:
 - Member Name
The name of the department or departments that are associated with the -memberid parameter of the department.
 - Member Type
The type of department or application.
- l
Specifies that the long version of the information is listed. Besides the **Member ID**, **Member Name** and **Member Type** columns, the following information is included:
 - UDP1, UDP2, UDP3
The user-defined property (UDP) contains anything of significance to the administrator. This alphanumeric character can be a maximum of 255 characters in length.

Example: Displaying a list of the members for a department

The following command generates a list of the members of a specified department:

```
tpctool> lsdeptgroupmembers -id/Department/5041
```

The following output is returned.

Member ID	Member Name	Member Type
/Department/5044	Tax	Department
/Application/6050	NetSuite	Application

Example: Displaying a long list of the members for a department

The following command generates a long list of the members of a specified department:

```
tpctool> lsdeptgroupmembers -id/Department/5041 -l
```

The following output is returned.

Member ID	Member Name	Member Type	UDP1	UDP2	UDP3
/Department/5044	Tax	Department	-	-	-
/Application/6050	NetSuite	Application	-	-	-


Related reference

- [mkdeptgroup](#)
- [modifydeptgroup](#)
- [rmdeptgroup](#)
- [lsdeptgroup](#)

Isdev

Use the **lsdev** command to list information about storage systems, fabrics, and switches. This information includes the globally unique identifier (GUID) or worldwide name (WWN) for the fabric, the user-defined name, the resource type, the status, and the time that the status was updated.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- perf
Specifies that resources for which performance data is collected will be listed. You must have the applicable authority to view the resources.
- level sample | hourly | daily
Specifies the level for which the information is summarized. You can specify a sample summary, an hourly summary, or a daily summary.
- fabric
Specifies that only fabrics are listed. You must have Administrator authority to use this option.
- sysys
Specifies that only storage systems are listed.
- l
Specifies that the long version of the information is listed. The following information is given.
 - GUID or WWN
 - User-defined name
 - Resource type
 - Status
 - Time that the status was updatedIf you omit this parameter, only the GUID or WWN is listed.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Command Aliases

You can run the following **lsdev** command aliases that are predefined by IBM®.

```
lsfabric
lsdev -fabric
lsperf
lsdev -perf
lssubsys
lsdev -subsys
```

Example: Listing all resources

The following command lists the GUIDs for all storage systems in the IBM Spectrum Control environment:

```
tpctool> lsdev -user me -pwd mypass -url myhost:myport -subsys
```

The following output is returned:

```
GUID
=====
9.47.97.159:0000020065400048+0
9.47.97.161:0000020060C0002A+0
1750.13AAW2A+0
1750.13AB1WA+0
2107.1302541+0
2107.1301901+0
2105.22232+0
2105.20870+0
```

Example: Listing the long version of information

The following command lists the long version of information for fabrics.

```
tpctool> lsdev -user me -pwd mypass -url myhost:myport -fabric -l
```

The following output is returned:

GUID	Name	Type	Status	Timestamp
100000051E34F6A8	100000051E34F6A8	-	UNKNOWN	2004.12.31:00:00:00
10000060695130FD	1000006069514262	-	UNKNOWN	2004.12.31:00:00:00
1000006069514262	10000060695130FD	-	UNKNOWN	2004.12.31:00:00:00

lsdevp

Use the **lsdevp** command to list worldwide port names (WWPNs) for a subsystem.

Syntax

```
tpctool lsdevp -user user_name -pwd password -url url -dev subsystem [-l] [-fs character]
[-header] [-help] [-qs character] [-?] [-silent]
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- dev subsystem
Specifies the globally unique identifier (GUID) of the storage subsystem. You can use the **lsdev** command to return information, including the GUID, for all storage subsystems that are discovered by IBM Spectrum Control.
- l
Specifies that the long version of the information is to be listed.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent

Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing worldwide port names

The following command lists the WWPNS for the specified subsystem:

```
tpctool> lsdevp -user me -pwd mypass -url myhost:myport -dev 2105.22232+0
```


The following output is returned:

```
WWPN
=====
5005076300C79470
5005076300D09470
5005076300CB9470
5005076300CC9470
5005076300C29470
5005076300CF9470
5005076300C89470
5005076300C39470
```

lsdisk

Use the **lsdisk** command to list all the physical disks on a specified storage subsystem.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- subsys subsystem
Specifies the GUID of the storage subsystem. You can use the **lsdev** command to return information, including the GUID, for all storage subsystems that are discovered.
- l
Specifies that long version of the information is listed. The following information is included.
 - Key
 - Label
 - Vendor
 - Model
 - Serial Number
 - Capacity
 - Speed
 - Encrypted
 - Solid State
- If this parameter is not issued, only the host name is returned.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Displaying a list of physical disks on a storage subsystem

The following command generates a list of physical disks on a specified storage subsystem:

```
lsdisk -subsys 2107-1300361+0 -l
```

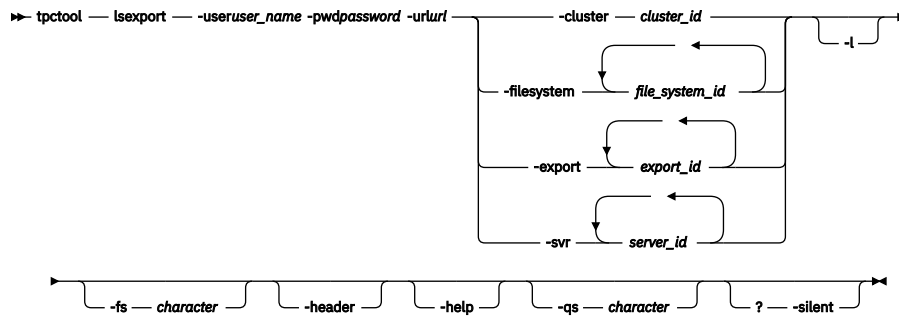
The following output is returned:

Key	Label			Vendor	Model
=====					
8000350BFC0D00D+IBM.2107-1300361	U2107.D01.Q000004-P1-D14			Seagate	S0AE146
Serial Number	Capacity	Speed	Encrypted	Solid State	
=====					
8000350BFC0D00D	146.0	-	No	No	

lsexport

Use the **lsexport** command to list all exports or specified exports that are associated with a cluster that is on a Storwize® V7000 Unified storage system. You can also list exports by file system, Netapp, or server that is managed by a Storage Resource Agent (SRA).

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- cluster cluster_id | -filesystem file_system_id... | -export export_id... | -svr svr_id...
The -cluster parameter specifies that information is displayed for all exports that are associated with a cluster. You can use the **lscluster** command to return information, including the cluster IDs, for all clusters that were discovered.
The -filesystem parameter specifies that information is displayed only for exports that are associated with specific file systems. You can enter one or multiple file system IDs. Use a comma to separate multiple IDs. You can use the **lsfs** command to return information, including the file system IDs, for all file systems that were discovered.
The -export parameter specifies that information is displayed only for specific exports. You can enter one or multiple export IDs. Use a comma to separate multiple IDs.
The -svr parameter specifies that information is displayed for all exports that are associated with a specific server. You can use the **lssvr** command to return information, including the server IDs, for all servers that were discovered.
- l
Specifies that the following information is listed.

ID
An ID for the export. For a file, the ID is composed of the export name, cluster name, storage system name, and storage system format. For a server, it is composed of the export name and server name.

Export
The name of the export.

Path
The path on the computer where the export is located.

Protocol
The protocol for the export can be one of following values: HTTP, FTP, SCP, CIFS, or NFS.

Active
Indicates whether the export is active or inactive in the cluster. The value can be true or false. If an export is inactive, it is included in the list of exports, but the data that is associated with the export cannot be accessed.

Options
The configuration options that were set for the export. This column contains information only if the protocol is CIFS or NFS.

If you omit this parameter, only the ID is returned.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent

Suppresses all output for the command. If you omit this parameter, output is included.

Example: List all nodes that are on a cluster

The following command generates a list of all exports that are on a Storwize V7000 Unified cluster.

```
tpctool> lsexport -cluster storage2.storage.tucson.ibm.com+00000200A0E0005C+0 -l
```

The following output is returned:

ID	Export
RandomExport1+storage2.storage.tucson.ibm.com+00000200A0E0005C+0	RandomExport1
RandomExport1+storage2.storage.tucson.ibm.com+00000200A0E0005C+0	RandomExport1
RandomExport1+storage2.storage.tucson.ibm.com+00000200A0E0005C+0	RandomExport1
testadd+storage2.storage.tucson.ibm.com+00000200A0E0005C+0	testadd
testadd+storage2.storage.tucson.ibm.com+00000200A0E0005C+0	testadd
smcho2+storage2.storage.tucson.ibm.com+00000200A0E0005C+0	smcho2
smcho2+storage2.storage.tucson.ibm.com+00000200A0E0005C+0	smcho2
smcho2+storage2.storage.tucson.ibm.com+00000200A0E0005C+0	smcho2
smcho2+storage2.storage.tucson.ibm.com+00000200A0E0005C+0	smcho2

Path	Active	Protocol	Options
/ibm/gpfs0/randomexport1	false	NFS	tb1435 (rw,no_wdelay)
/ibm/gpfs0/randomexport1	false	CIFS	access control=Everyone:ALLOWED:FULL
/ibm/gpfs0/randomexport1	false	HTTP	
/ibm/testadd	true	NFS	tb1435 (rw,no_wdelay)
/ibm/testadd	true	CIFS	read only;access control=Everyone:

ALLOWED:FULL			
/ibm/gpfs0/smcho	true	CIFS	access control=Everyone:ALLOWED:FULL
/ibm/gpfs0/smcho	true	HTTP	
/ibm/gpfs0/smcho	true	FTP	
/ibm/gpfs0/smcho	true	SCP	

lsextent

Use the **lsextent** command to display a list of all the storage extents on a specified storage system. An example of a storage extent is a managed disk on a SAN Volume Controller. For Dell EMC VMAX storage systems, the command displays the Disk Group.

Syntax

```
tpctool — lsextent — -user user_name -pwd password -url url — -sys sys — subsystem
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- sys subsystem
Specifies the GUID of the storage system. You can use the **lsdev** command to return information, including the GUID, for all storage systems that are discovered.
- l
Specifies that long version of the information is listed. The following information is included.
 - Key
 - Label
 - Total Capacity
 - Free Space
 - Virtual
- If this parameter is not issued, only the host name is returned.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Displaying a list of storage extents on a storage system

The following command generates a list of storage extents on a specified storage system:

```
tpctool> lsextent -subsys 2107-1300361+0 -l
```

The following output is returned:

Key	Label	Total Capacity
IBM.2107-1300361-R8+2107-1300361+0	IBM.2107-1300361-R8	581.0
IBM.2107-1300361-R7+2107-1300361+0	IBM.2107-1300361-R7	1688.0
IBM.2107-1300361-R6+2107-1300361+0	IBM.2107-1300361-R6	519.0

Free Space Virtual
0.0 No
0.0 No
0.0 No

lsfcpath

Use the **lsfcpath** command to list the paths for data transmission between a system with a fibre-channel host bus adapter (HBA) and a storage subsystem.

Syntax

```
tpctool lsfcpath -user user_name -pwd password -url url -svr server -dev GUID -fs character -header -help -qs character ? -silent
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- svr server
Specifies a system with a fibre-channel HBA. The *server* variable is the host name of the system.
- dev GUID
Specifies the storage subsystem. The *GUID* variable is the globally unique identifier (GUID).
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing data paths

The following command lists the paths on which data can be transmitted between MARKETING and the storage subsystem 2105.20870+0:

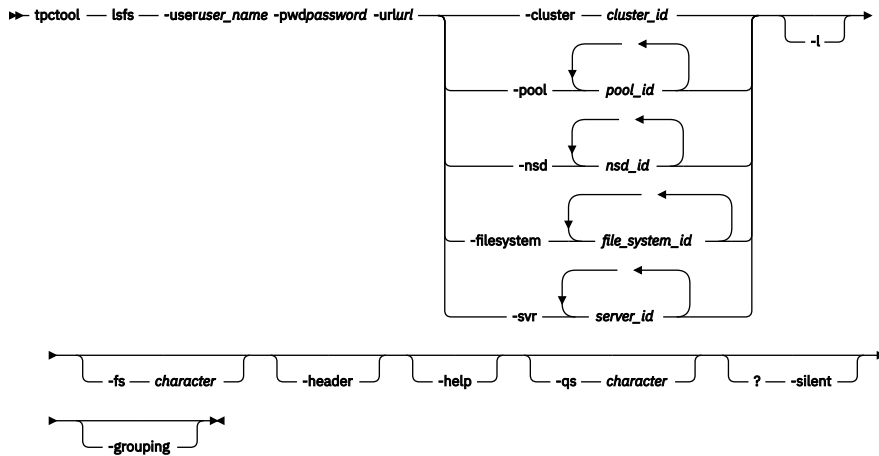
```
tpctool> lsfcpath -user me -pwd mypass -url myhost:myport -svr MARKETING -dev 2105.20870+0
```

The following output is returned:

ServerPort	SubsystemPort
210000E08B1Co9E	710000E08W1Co8F

lsfs

Use the **lsfs** command to list all file systems or specified file systems that are associated with a cluster on a Storwize® V7000 Unified or IBM Spectrum Scale storage system. You can also list file systems by pool or Network Shared Disk (NSD).



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- cluster cluster_id | -pool pool_id... | -nsd nsd_id... | -filesystem file_system_id... | -svr server_id...
The -cluster parameter specifies that information is displayed for all file systems that are on a cluster. You can use the **lscluster** command to return information, including the cluster IDs, for all clusters that were discovered.
The -pool parameter specifies that information is displayed only for file systems that are associated with specific pools. You can enter one or multiple pool IDs. Use a comma to separate multiple IDs. You can use the **lspool** command to return information, including the pool IDs, for all pools that were discovered.
The -nsd parameter specifies that information is displayed only for file systems that are associated with specific NSDs. You can enter one or multiple NSD IDs. Use a comma to separate multiple IDs. You can use the **lsnsd** command to return information, including the NSD IDs, for all NSDs that were discovered.
The -filesystem parameter specifies that information is displayed only for specific file systems. You can enter one or multiple file system IDs. Use a comma to separate multiple IDs.
The -svr parameter specifies that information is displayed for all file systems that are associated with a specific server. You can use the **lssvr** command to return information, including the server IDs, for all servers that were discovered.
- l
Specifies that the following information is listed.
 - ID
fqz0_r_cli_lsfs
fqz0_r_cli_lsfs
An ID for the file system that is composed of the file system name, cluster name, storage system name, and storage system format.
 - File System
The name of the file system.
 - Mount Point
The name or mount point (UNIX/Linux®) of the file system (for example, c:\, d:\, /opt, or /export/home).
 - Cluster/Server
The name of the cluster or server that is associated with the file system.
 - Capacity (GB)
The capacity of the file system.
 - Used Space (GB)
The amount of used storage space in the file system.
 - Available Space (GB)
The amount of unused storage space in the file system.
 - Maximum File Count
The total number of files in a file system. This value does not include files on file systems that were not scanned.
 - Used I-Nodes
The number of used inodes in the file system.
 - Free I-Nodes
The number of unused inodes in the file system. If you omit this parameter, only the ID is returned.
 - UDP1
The user-defined property (UDP) contains anything of significance to the group administrator. This alphanumeric character string must contain fewer than 255 characters in length.
 - UDP2
The user-defined property (UDP) contains anything of significance to the group administrator. This alphanumeric character string must contain fewer than 255 characters in length.
 - UDP3
The user-defined property (UDP) contains anything of significance to the group administrator. This alphanumeric character string must contain fewer than 255 characters in length.

- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- grouping
Enables the grouping of numeric values. For example, in English, the value 12000 displays as 12,000. The grouping character is determined by the system locale.

Example: List all file systems that are on a cluster

The following command generates a list of all file systems that are on a Storwize V7000 Unified cluster.

```
tpctool> lsfs -cluster storage2.storage.abc.ibm.com+00000200A0E0005C+0 -l
```

The following output is returned:

ID	File System
gpfs0+storage2.storage.abc.ibm.com+00000200A0E0005C+0	gpfs0
testadd+storage2.storage.abc.ibm.com+00000200A0E0005C+0	testadd
gpfs1+storage2.storage.abc.ibm.com+00000200A0E0005C+0	gpfs1
ma_19037+storage2.storage.abc.ibm.com+00000200A0E0005C+0	ma_19037

Cluster	Mount Point	Capacity (GB)	Used Space (GB)
storage2.storage.abc.ibm.com	/ibm/gpfs0	4280	1135.91
storage2.storage.abc.ibm.com	/ibm/testadd	8560	151.58
storage2.storage.abc.ibm.com	/ibm/scantest	4280	1.54
storage2.storage.abc.ibm.com	/ibm/gpfs0/ma	4280	1.54


Available Space (GB)	Maximum File Count	Used I-Nodes	Free I-Nodes
3144.09	4382726	537458	3845268
8408.42	4383232	5107	4378125
4278.46	4382726	4044	4378682
4278.46	4382726	4043	4378683

UDP1	UDP2	UDP3
GPFS	root	-
GPFS	nonroot	-
MA	nonroot	19037FS

lsfs

Use the **lsfs** command to list all filesets or specified filesets that are associated with a file system on a Storwize® V7000 Unified or IBM Spectrum Scale storage system.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- filesystem file_system_id | -fileset file_set_id...
Specifies that information is displayed for all filesets that are on a file system.
The -filesystem parameter requires the ID for the file system. You can use the **lsfs** command to view information, including the file system IDs, for all file systems that were discovered.

The -fileset parameter specifies that information is displayed only for specific filesets. You can enter one or multiple fileset IDs. Use a comma to separate multiple IDs.

-l

Specifies that information is listed. The following information is included:

ID

An ID for the fileset that is composed of the fileset name, file system name, cluster name, storage system name, and storage system format.

Fileset

The name of the fileset.

Path

The path for the fileset. The path is displayed only if linked is displayed in the Status column.

Status

The status of the fileset can be one of the following values: linked or unlinked.

Used Space (GB)

The amount of space that is used by the fileset.

Used I-Nodes

The number of used inodes in the fileset.

Type

The type of the fileset, such as cache source or cache. No value in this column indicates that the fileset is not enabled for caching.

If you omit this parameter, only the ID is returned.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

-grouping

Enables the grouping of numeric values. For example, in English, the value 12000 is displayed as 12,000. The grouping character is determined by the system locale.

Example: List all filesets that are on a specified file system

The following command generates a list of all filesets that are on a file system:

```
tpctool> lsfset -filesystem  
gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0  
-l
```

The following output is returned:

ID

```
=====
root+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
CiprianCet+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
p_swfset1+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
afs_c_ro3+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
afs_c_ro43+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
afs_c_ro5+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
cachefset1+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
Cet1+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
testCache+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
testCache_rw+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
eecache01+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
shirley_cache_source_fset+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
afs_c+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
adriand_ace+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
afs_c_ro+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
afs_c_lu+gpfs0+storage3.storage.tucson.ibm.com+127.0.0.1+0
```

Fileset	Path	Status	Used Space (GB)
root	/ibm/gpfs0	linked	3.8
CiprianCet	/ibm/gpfs0/CiprianCet	linked	0
p_swfset1	/ibm/gpfs0/p_swfset1	linked	-
afs_c_ro3	/ibm/gpfs0/afs_c_ro3	linked	-
afs_c_ro43	/ibm/gpfs0/afs_c_ro4	linked	-
afs_c_ro5	/ibm/gpfs0/afs_c_ro5	linked	-
cachefset1	/ibm/gpfs0/cachefset1	linked	-
Cet1	/ibm/gpfs0/CiprianCet/Cet1	linked	0
testCache	/ibm/gpfs0/testCache	linked	-
testCache_rw	/ibm/gpfs0/testCache_rw	linked	-
eecache01	/ibm/gpfs0/eecache01	linked	-
shirley_cache_source_fset	/ibm/gpfs0/shirley_cache_source_fset	linked	-
afs_c	/ibm/gpfs0/afs_c	linked	-
adriand_ace	/ibm/gpfs0/adriand_acedir	linked	0
afs_c_ro	/ibm/gpfs0/afs_c_ro	linked	-
afs_c_lu	/ibm/gpfs0/afs_c_lu	linked	-
Used I-Nodes	Type		
4110	-		
4	Cache source		

4	Cache
3	Cache
2	Cache
7	Cache
2	Cache
1	-
3	Cache
3	Cache
2	Cache
2	Cache
2	Cache
1	-
2	Cache
2	Cache


Related reference

- [lsfs](#)

Ishtype

Use the **Ishtype** command to list host types. You must have Administrator authority to use this command.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- dev subsystem
Specifies the globally unique identifier (GUID) of the storage subsystem. You can use the **lsdev** command to obtain information, including the GUID, for all storage subsystems that are discovered.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing host types

The following command lists the host types that are associated with the specified subsystem.

```
tpctool> lshtype -user me -pwd mypass -url myhost:myport -dev 2105.22232+0
```

The following output is returned.

```
HostType HostFlag
=====
StorageClientSettingData for IBM pSeries, RS/6000 and RS/6000 SP Servers (AIX) 9
StorageClientSettingData for Cisco iSCSI Gateway 13
StorageClientSettingData for DGUX 10
StorageClientSettingData for IBM NUMA-Q Servers (DYNIX/ptx) 11
StorageClientSettingData for HP Servers (HP-UX) 4
StorageClientSettingData for SGI Origin Servers (IRIX) 12
StorageClientSettingData for Intel-based Servers (Linux) 14
StorageClientSettingData for IBM zSeries Servers (Linux) 14
StorageClientSettingData for IBM iSeries/pSeries Servers (Linux) 14
StorageClientSettingData for Intel-based Servers (Microsoft Windows NT4) 15
StorageClientSettingData for Intel-based Servers (Novell NetWare) 7
StorageClientSettingData for HP AlphaServer (OpenVMS) 5
StorageClientSettingData for IBM AS/400 (V3R7 to V5R2) 16
StorageClientSettingData for IBM AS/400 (V5R3 or higher) 16
StorageClientSettingData for IBM SAN File System (AIX MDS) 9
```

```
StorageClientSettingData for IBM SAN File System (Linux MDS) 14
StorageClientSettingData for IBM SAN Volume Controller 3277 0
StorageClientSettingData for Solaris 2.51 3
StorageClientSettingData for Sun Servers (Solaris 2.6 or higher) 3
StorageClientSettingData for Sun Clustering using MPxIO 3279 0
StorageClientSettingData for HP AlphaServer (Tru64 UNIX) 6
StorageClientSettingData for VMware ESX 3278 0
tpctool>
```

Isoptschedules

Use the **isoptschedules** command to show a list of the schedules that you created to analyze storage tiering.

In the GUI, you can analyze storage tiering. On the Optimization Analysis page, you can create schedules that are based on the criteria that you created when you analyzed tiering.

When you issue the **isoptschedules** command, the following information is shown:

Schedule ID

Specifies the ID of the schedule.


Schedule Name

Specifies the name of the schedule.

Most Recent Job ID

Specifies the current or most recent job ID.

Syntax



Parameters and arguments

-l

Specifies that the following additional information is listed.

Most Recent Start Time

The last time that the schedule was run.

Next Start Time

The next time that the schedule is run.

-user user_name

Specifies an IBM Spectrum Control user ID.

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

Example: Showing a list of the analyze tiering schedules

Issue the following command to show a list of the schedules that were created to analyze tiering:

```
tpctool> isoptschedules -l
```

The following output is returned:


Schedule ID	Schedule Name	Most Recent Job ID
1244	opt1357942625952	10002
459011	op61354	8038

Most Recent Start Time	Next Start Time
2013.01.11:15:11:05	2013.02.11:17:00:30
2013.01.12:07:33:12	NA

Ismetrics

Use the **lsmetrics** command to list available performance metrics. You must have Fabric operator or Disk operator authority to use this command.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- subsys subsystem
Specifies the storage subsystem. The subsystem variable is the GUID of the storage subsystem. You can use the **lsdev** command to return information, including the GUID, for all storage subsystems that are discovered by IBM Spectrum® Control.
- ctype comp_type
Specifies that the output includes only components of the specified type. For more information about the *comp_type* variable, see the **lstype** command.
- level sample | hourly | daily
Specifies the level for which the performance metrics be summarized. You can specify a sample summary, an hourly summary, or a daily summary.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing performance metrics

The following command generates a list of performance metrics for switch ports on the specified fabric:

```
tptool lsmetrics -user me -pwd mypass -url myhost:myport  
-fabric 100000051E34F6A8 -ctype switch_port -level sample
```

The following output is returned:

Metric	Value
Port Send Packet Rate	855
Port Receive Packet Rate	856
Total Port Packet Rate	857
Port Send Data Rate	858
Port Receive Data Rate	859
Total Port Data Rate	860
Port Peak Send Data Rate	861
Port Peak Receive Data Rate	862
Port Send Packet Size	869
Port Receive Packet Size	870
Overall Port Packet Size	871
Error Frame Rate	872
Dumped Frame Rate	873
Link Failure Rate	874


Related reference

- [lstype](#)

Isnode

Use the **Isnode** command to list all nodes or specified nodes that are associated with a cluster that is on a Storwize® V7000 Unified storage system.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- cluster cluster_id | -node node_id...
Specifies that information is displayed for all nodes that are on a cluster.
For the -cluster parameter, specify the ID for the cluster. You can use the **lscluster** command to view information, including the cluster IDs, for all clusters that were discovered.
The -node parameter specifies that information is displayed only for specific nodes. You can enter one or multiple node IDs. Use a comma to separate multiple IDs.
- l
Specifies that the following information is listed.

ID
An ID for the node that is composed of the node name, cluster name, storage system name, and storage system format.

Node
The name of the node.

Description
The description of the node.

IP Address
The IP address for the node.

Role
One or more of the following roles: interface, management, or storage.

Cache Gateway Node
Specifies **Yes** or **No** to indicate whether an interface node is enabled to function as a caching gateway node that exchanges data with other systems.

Status
One of the following values: **NORMAL**, **CRITICAL**, or **WARNING**.

If you omit this parameter, only the ID is returned.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing all nodes that are on a cluster

The following command generates a list of all nodes that are on a cluster.

```
tpctool> lsnode -cluster storage3.storage.tucson.ibm.com+127.0.0.1+0 -l
```

The following output is returned:

```
ID
=====
int001st001+storage3.storage.tucson.ibm.com+127.0.0.1+0
int002st001+storage3.storage.tucson.ibm.com+127.0.0.1+0
int003st001+storage3.storage.tucson.ibm.com+127.0.0.1+0
mgmt001st001+storage3.storage.tucson.ibm.com+127.0.0.1+0
strg001st001+storage3.storage.tucson.ibm.com+127.0.0.1+0
strg002st001+storage3.storage.tucson.ibm.com+127.0.0.1+0

Node      Description  IP Address  Role
=====
int001st001  int001st001  127.0.0.1  interface
int002st001  int002st001  127.0.0.1  interface
int003st001  int003st001  127.0.0.1  interface
mgmt001st001 mgmt001st001  127.0.0.1  interface,management
strg001st001 strg001st001  127.0.0.1  storage
strg002st001 strg002st001  127.0.0.1  storage
```

Status	Cache Gateway Node
NORMAL	Yes
NORMAL	Yes
NORMAL	No
NORMAL	No
NORMAL	No
NORMAL	No

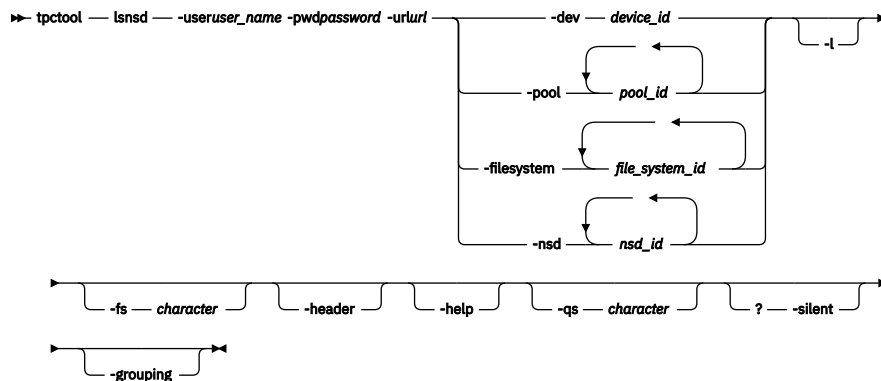
Related reference

- [lscluster](#)

lsnsd

Use the **lsnsd** command to list all Network Shared Disks (NSDs) or specified NSDs that are on a Storwize® V7000 Unified storage system. You can also list NSDs by file system, pool, or both.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- dev dev_id | -pool pool_id... | -filesystem file_system_id... | -nsd nsd_id...
The -dev parameter specifies that information is displayed for all NSDs that are associated with a storage system. This parameter requires the globally unique identifier (GUID) of the storage system. You can use the **lsdev** command to return information, including the GUID, for all storage systems that were discovered.
The -pool parameter specifies that information is displayed only for NSDs that are associated with specific pools. You can enter one or multiple pool IDs. Use a comma to separate multiple IDs. You can use the **lspool** command to return information, including the pool IDs, for all pools that were discovered.
The -filesystem parameter specifies that information is displayed only for NSDs that are associated with specific file systems. You can enter one or multiple file system IDs. Use a comma to separate multiple IDs. You can use the **lsfs** command to return information, including the file system IDs, for all file systems that were discovered.
The -nsd parameter specifies that information is displayed only for specific NSDs. You can enter one or multiple NSD IDs. Use a comma to separate multiple IDs.
- l
Specifies that the following information is listed.
 - ID
The ID for the NSD that is composed of the NSD name, storage system name, and storage system format.
 - NSD
NSD: The name of the NSD.
 - Cluster ID
The ID for the cluster that is associated with the NSD as defined by the Storwize V7000 Unified file module.
 - Failover Group
The failure group that the NSD belongs to.
 - Type
One or more of the following types: dataAndMetadata, dataOnly, metadataOnly.
 - Pool
The pool that the NSD is associated with.
 - Status
One of the following values: NORMAL, CRITICAL, WARNING, or UNREACHABLE.
 - Disk Space (GB)
The total disk space for the NSD.
 - Available Space (GB)
The amount of unused disk space for the NSD.

File System

The file system that the NSD is associated with.

If you omit this parameter, only the ID is returned.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

-grouping

Enables the grouping of numeric values. For example, in English, the value 12000 displays as 12,000. The grouping character is determined by the system locale.

Example: List all NSDs that are on a storage system

The following command generates a list of all NSDs that are on a Storwize V7000 Unified storage system.

```
tpctool> lsnsd -dev 00000200A0E0005C+0 -1
```

The following output is returned:

ID		NSD	
array0_sas_60001ff07996c0089b00000+00000200A0E0005C+0		array0_sas_60001ff07996c0089b00000	
array0_sas_60001ff07996c0289b20002+00000200A0E0005C+0		array0_sas_60001ff07996c0289b20002	
array1_sas_60001ff07996c0389b30003+00000200A0E0005C+0		array1_sas_60001ff07996c0389b30003	
array1_sas_60001ff07996c0189b10001+00000200A0E0005C+0		array1_sas_60001ff07996c0189b10001	
array0_sas_60001ff07996c0489b40004+00000200A0E0005C+0		array0_sas_60001ff07996c0489b40004	

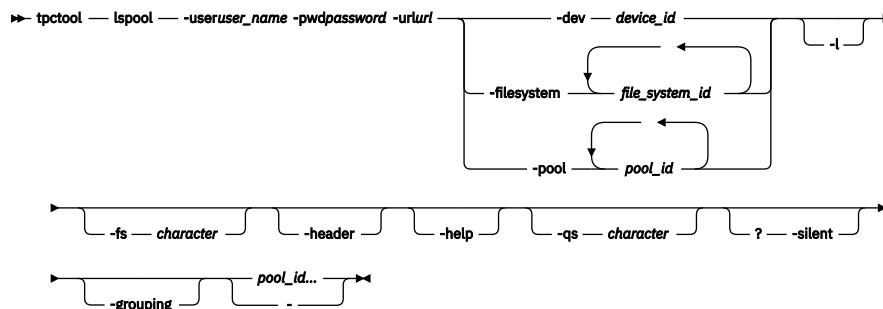
Cluster ID	Failover Group	Type	Pool	Status
12402779238946656959	1	dataAndMetadata	system	NORMAL
12402779238946656959	1	dataAndMetadata	system	NORMAL
12402779238946656959	4005	dataAndMetadata	system	NORMAL
12402779238946656959	1	dataAndMetadata	system	NORMAL
12402779238946656959	4002	dataAndMetadata	system	NORMAL

Disk Space (GB)	Available Space (GB)	File System
4280	3144	gpfs0
4280	4278.46	gpfs1
4280	4278.46	ma_19037
4280	4128.4	testadd
4280	4280	testadd

lspool

Use the **lspool** command to list all file system pools that are on a specified Storwize® V7000 Unified storage system. You can also list the pools by file system.

Syntax



Parameters and arguments

-user user_name

Specifies an IBM Spectrum Control user ID.

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-dev dev_id | -filesystem file_system_id... | -pool pool_id...

The `-dev` parameter specifies that information is displayed for all pools that are associated with a storage system. This parameter requires the globally unique identifier (GUID) of the storage system. You can use the `lsdev` command to return information, including the GUID, for all storage systems that were discovered. The `-filesystem` parameter specifies that information is displayed only for pools that are associated with specific file systems. You can enter one or multiple file system IDs. Use a comma to separate multiple IDs. You can use the `lsfs` command to return information, including the file system IDs, for all file systems that were discovered. The `-pool` parameter specifies that information is displayed only for specific pools. You can enter one or multiple pool IDs. Use a comma to separate multiple IDs.

- l
Specifies that the following information is listed.
 - ID
An ID for the pool that is composed of the pool name, storage system name, and storage system format.
 - Pool
The name of the pool.
 - Capacity (GB)
The capacity of the pool.
 - Available Space (GB)
The amount of unused space that is in the pool.
- If you omit this parameter, only the ID is returned.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- grouping
Enables the grouping of numeric values. For example, in English, the value 12000 displays as 12,000. The grouping character is determined by the system locale.

Example: List all file system pools that are on a storage system

The following command generates a list of all file system pools that are on a Storwize V7000 Unified storage system.

```
tpctool> lspool -dev 00000200A0E0005C+0 -l
```

The following output is returned:

ID	Pool	Capacity (GB)	Available Space (GB)
system+00000200A0E0005C+0	system	21400	20109.43

lsport

Use the `lsport` command to list the ports that are on a Fibre Channel host bus adapter (HBA).

Syntax

```
tpctool lsport -user user_name -pwd password -url url -svr server [-l] [-fs character]
[-header] [-help] [-qs character] [-?] [-silent]
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is `system:port_number`, where `system` represents either the host name or IP address, and `port_number` represents the IBM Spectrum Control Device server port.
- svr server
Specifies a system with a fibre-channel HBA. The `server` variable is the host name of the system.
- l
Specifies that the long version of the information is listed.
 - Worldwide port name (WWPN)
 - Name
 - Status
- If you omit this parameter, only the WWPN is returned.
- fs character

- Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
 - Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
 - Lists help information for the command.
- qs character
 - Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
 - Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing port information

The following command lists the long version of information about the ports that are associated with the Fibre Channel HBA.

```
tpctool> lsport -user me -pwd mypass -url myhost:myport -svr <server> -l
```

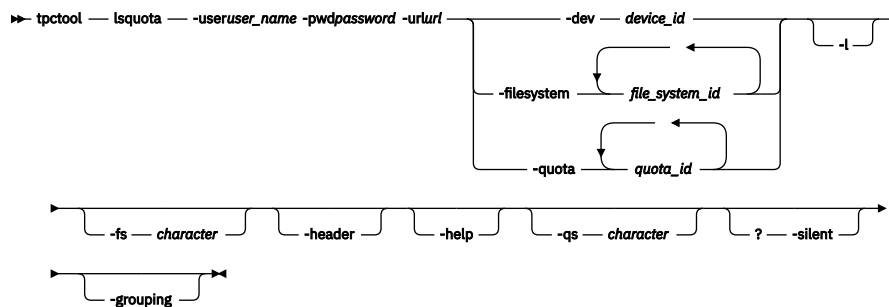
The following output is returned:

WWPN	Name	Status
210000E08B1Co9E	Marketing	Active

lsquota

Use the **lsquota** command to list all quotas or specified quotas that are on a Storwize® V7000 Unified storage system. You can also list quotas by file system.

Syntax



Parameters and arguments

- user user_name
 - Specifies an IBM Spectrum Control user ID.
 - pwd password
 - Specifies the password for the IBM Spectrum Control user ID.
 - url url
 - Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
 - dev dev_id | -filesystem file_system_id... | -quota quota_id...
 - The -dev parameter specifies that information is displayed for all quotas that are associated with a storage system. This parameter requires the globally unique identifier (GUID) of the storage system. You can use the **lsdev** command to return information, including the GUID, for all storage systems that were discovered.
 - The -filesystem parameter specifies that information is displayed only for quotas that are associated with specific file systems. You can enter one or multiple file system IDs. Use a comma to separate multiple IDs. You can use the **lsfs** command to return information, including the file system IDs, for all file systems that were discovered.
 - The -quota parameter specifies that information is displayed only for specific quotas. You can enter one or multiple quota IDs. Use a comma to separate multiple IDs.
 - l
 - Specifies that the following information is listed.
- | | |
|-----------------|---|
| ID | The identifier for the quota. |
| Name | The name of the user, group, or fileset that the quota is associated with. |
| File System | The file system that the quota is associated with. |
| Type | The type of quota can be one of the following values: user, file set, or group. |
| SL Usage (GB) | The soft limit for storage space usage. |
| HL Usage (GB) | The hard limit for storage space usage. |
| Used Space (GB) | |

- The amount of used storage space.
 - SL I-Nodes
 - The soft limit for inode usage.
 - HL I-Nodes
 - The hard limit for inode usage.
 - Used I-Nodes
 - The number of used inodes.
 - Gracetime Usage (sec.)
 - The time frame in which storage space usage must be brought below the quota.
 - Gracetime I-Nodes (sec.)
 - The time frame in which inode usage must be brought below the quota.
 - In Doubt (KB)
 - The amount of data for which the quota was not updated.
- If you omit this parameter, only the ID is returned.
- fs character
 - Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
 - header
 - Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
 - help | -h | -?
 - Lists help information for the command.
 - qs character
 - Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
 - silent
 - Suppresses all output for the command. If you omit this parameter, output is included.
 - grouping
 - Enables the grouping of numeric values. For example, in English, the value 12000 displays as 12,000. The grouping character is determined by the system locale.

Example: Displaying a list of all quotas that are on a storage system.

The following command generates a list of all quotas on a Storwize V7000 Unified storage system:

```
tpctool> lsquota -dev 00000200A0E0005C+0 -l
```

The following output is returned:

ID	Name						
4572_U_10000000+00000200A0E0005C+0	10000000						
4572_F_Adrian+00000200A0E0005C+0	Adrian						
4578_U_SRM\administrator+00000200A0E0005C+0	SRM\administrator						
4574_U_SRM\administrator+00000200A0E0005C+0	SRM\administrator						
4572_U_SRM\administrator+00000200A0E0005C+0	SRM\administrator						

File System	Type	SL Usage (GB)	HL Usage (GB)	Used Space (GB)				
gpfs0			user	0	0	0	0	
gpfs0			fileset	0	0	0	0	
testadd		user	0	0	0	0		
gpfs1		user	0	0	0	0.02		
gpfs0			user	0	0	579654.05		

SL I-Nodes	HL I-Nodes	Used I-Nodes	
0	0	1	
0	0	1	
512	512	0	
0	0	4	
0	0	533268	

Gracetime Usage (sec.)	Gracetime I-Nodes (sec.)	In Doubt
0	0	0
0	0	0
0	0	0
0	0	0
0	0	0

lssrg

Use the **lssrg** command to display a list of storage resource groups that are known to IBM Spectrum Control.

Syntax

```
tpctool — lssrg — -user user_name -pwd password -url url
                                     -fs character -header -help
                                     -qs character ? -silent -l
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
 - pwd password
Specifies the password for the IBM Spectrum Control user ID.
 - url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
 - fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
 - header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
 - help | -h | -?
Lists help information for the command.
 - qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
 - silent
Suppresses all output for the command. If you omit this parameter, output is included.
 - l
Specifies that the long version of the information is listed, which includes the following information:
 - Description
 - User Defined Property 1
 - User Defined Property 2
 - User Defined Property 3
- If this parameter is not issued, only the group name is returned.

Example: Displaying a list of storage resource groups

The following command generates a list of storage resource groups that are known to IBM Spectrum Control.

```
tpctool>lssrg -l
```

The following output is returned:

Name	Description	UDP1	UDP2	UDP3
=====				
administrator.my-example-srg				

lssrgmembers

Use the **lssrgmembers** command to list members of a specified storage resource group.

Syntax

```
tpctool -- lssrgmembers -- -user user_name -pwd password -url url  
                                     -fs character -header  
-help -qs character ? -silent -name SRGName
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- name SRGName
The unique name of a storage resource group. The output includes the following information.

- Member Key
- Member Type

Example: Displaying a list of the members of a storage resource group

The following command generates a list of the members of a specified storage resource group.

```
tpctool> lssrgmembers -name administrator.my-example-srg
```

The following output is returned.

Member Key	Member Type
1000080088E32D2D	Switch
100000051E0405C5	Switch
2107-1300361+0	Subsystem

lssvr

Use the **lssvr** command to list all systems that are discovered by Fabric Manager. You must have Fabric Administrator authority to use this command.

Syntax

```
tpctool -- lssvr -- -user user_name -pwd password -url url [-l] [-help] [-fs character]
[-qs character] [-? -silent] [-header]
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- l
Specifies that the long version of the information is listed. The following information is included:
 - Host name
 - Operating system
 - IP address
 - Status
 - UDP1, which can contain any alphanumeric character, and must be fewer than 255 characters in length.
 - UDP2, which can contain any alphanumeric character, and must be fewer than 255 characters in length.
 - UDP3, which can contain any alphanumeric character, and must be fewer than 255 characters in length.
 - Key - a combination of the GUID, the host name, and the IP address

If this parameter is not issued, only the host name is returned.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing long version of the information for all systems discovered by Fabric Manager

The following commands list the long version of the information for all systems that are discovered by Fabric Manager.

```
tpctool
tpctool> lssvr -l
```

The following output is returned:

Name	OS	IP	Status
Marketing	Windows	9.32.245.164	NORMAL


Key
bfcf0e0192a111dc9ac60011258d9a29+tb142-wi+9.47.97.149+++

UDP1	UDP2	UDP3
WindowsSRA	ABCLab	Bld80

lssvrdisk

Use the **lssvrdisk** command to list all of the physical disks that are known to a specified server.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- srv server_name
Specifies the key of a server that has a IBM Spectrum® Control agent that is deployed to it. Use the **lssvr -l** command to retrieve a list of servers and the associated keys.
- l
Specifies that long version of the information is listed. The following information is included.
 - Key - a combination of the GUID, the host name, and the IP address
 - Driver
 - Vendor
 - Model
 - Serial Number
 - Capacity
 - Path

If this parameter is not issued, only the host name is returned.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing long information for all physical disks known to a specified server

The following commands list long information for all physical disks that are known to a specified server.

```

tpctool
tpctool> lssvrdisk -srv bfcf0e0192a111dc9ac60011258d9a29+tb142-wi.beaverton.
ibm.com+9.47.97.149+++ -l

```

The following output is returned:


Key	Driver	Vendor	Model
3HX30J0X000075154506	Adaptec AIC-7902B - Ultra320 SCSI	IBM-ESXS	ST336753LC FN

Serial Number	Capacity	Path
3HX30J0X000075154506	36401479680	Disk 0

lsswitch

Use the **lsswitch** command to display a list of all switches in a specified fabric.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- fabric fabric
Specifies the WWN of the fabric that is associated with the switches. Use the **lsdev -l** command to see a list of available fabrics.
- l
Specifies that the long version of the information is listed. The following information is included.

- WWN
- Label
- Type
- Vendor
- Model
- Status

If this parameter is not issued, only the host name is returned.

Example: Listing long information for all switches on a specified fabric

The following command lists long information for all switches in the specified fabric:

```

tpctool
tpctool> lsswitch -fabric 10000800880476F0 -l

```


The following output is returned:

WWN	Label	Type	Vendor	Model	Status
100000051E34AEE1	tb500_sw	Physical	Unknown	26.2	UNREACHABLE

lstime

Use the **lstime** command to list the time ranges for which performance data is available.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url

- Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
 - subsys subsystem
Specifies the storage subsystem. The *subsystem* variable is the GUID of the storage subsystem. You can use the **lsdev** command to return information, including the GUID, for all storage subsystems that are discovered by IBM Spectrum® Control.
 - level sample | hourly | daily
Specifies the level for which the time ranges for performance data that is available is to be summarized. You can specify a sample summary, an hourly summary, or a daily summary.
 - ctype comp_type
Specifies that the output include only components of the specified type. For more information about the *comp_type* variable, see the **lstype** command.
 - comps component_list
Specifies the component list. The *component_list* variable specifies the components, such as those that are returned by **lscomp**.
 - fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
 - header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
 - help | -h | -?
Lists help information for the command.
 - qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
 - silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Displaying a list of time ranges

The following command generates a list of time ranges for performance data.

```
tpctool> lstime -user me -pwd mypass -url myhost:myport
-fabric 100000051E34F6A8 -ctype switch_port -level sample
```


Related reference

- [lstype](#)

lstype

Use the **lstype** command to list the components that are recognized by IBM Spectrum® Control. No authorization is required to run this command.

Syntax



Parameters and arguments

- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing all component types

The following command lists the component types that are recognized by IBM Spectrum Control:

```
tpctool> lstype
```

The following output is returned:

Name	Type
all	-1
unknown	0
subsystem	1

```


subsys_port 2
controller 3
stor_pool 4
svc_iogrp 5
ds_rio 6
svc_mdgrp 7
da 8
ds_rank 9
array 10
svc_mdisk 11
vol 12
switch 13
switch_port 14

```

lsvm

Use the **lsvm** command to list all virtual machines that are known to IBM Spectrum® Control. You can also list only the virtual machines that were discovered through hypervisors.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- hypervisor hypervisor_id
Specifies that information is displayed for all virtual machines that were discovered through a specific hypervisor.
- l
Specifies that the following information is listed.
 - VM ID
A unique ID for the virtual machine.
 - VM Name
The name of the virtual machine if discovered through a hypervisor. For all other virtual machines, the name is the URL of the host server.
 - Hypervisor Name
The name of the hypervisor for virtual machines that were discovered through a hypervisor. For all other virtual machines, this field is non-applicable.
 - IP Address
The IP address of the server that is associated with the virtual machine.
 - OS Type
The operating system type of the server that is associated with the virtual machine.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: List all virtual machines on a hypervisor

The following command generates a list of all virtual machines for a particular hypervisor.

```
tpctool> lsvm -l
```

The following output is returned:

```


ID Name Hypervisor Name OS Type Host Name IP Address
=====
sor1+sour.storage.abc.ib.com sor1 sour.storage.abc.ib.com "Microsoft Windows Server 2016 or later (64-Bit)"
sor1.storage.abc.ib.com 192.168.111.21
sor2+sour.storage.abc.ib.com sor2 sour.storage.abc.ib.com "Red Hat Enterprise Linux 7 (64-Bit)"
sor2.storage.abc.ib.com 192.168.111.22

```

lsvm disk

Use the **lsvm disk** command to list all the VMWare virtual machine disks that are known to the given virtual machine or hypervisor.

Syntax




Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- svr servername
Specifies a server key that represents either a virtual machine or hypervisor. Server keys can be obtained by using the **lssvr -l** command.
- l
Specifies that the long version of the information is listed. If this parameter is not issued, only the host name is returned.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

lsvol

Use the **lsvol** command to list all volumes on a system, list a specific volume or volumes, or list volumes on a specific array.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- dev GUID
Specifies the storage subsystem. The *GUID* variable is the globally unique identifier (GUID).
- array array
Specifies the array. The *array* variable is the array ID. You can use the **lsarray** command to return information, including array IDs, about the arrays on a specific storage subsystem.
- l
Specifies that long information is listed. In addition to the volume ID, the volume size and FlashCopy® relationship information are listed. Also includes the user defined properties: UDP1, UDP2 and UDP3.
The user-defined property (UDP) contains anything of significance to the group administrator. This alphanumeric character string must contain fewer than 255 characters in length.

- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- volume id | -
Specifies the volumes. The *volume_ID* variable is a comma-separated list of volume IDs, such as that obtained by running the **lsvol** command. If a single dash (-) is issued, the volume IDs are read from standard input.

Example: Listing the volumes on a subsystem

The following command lists all volumes on a specified subsystem. The **-dev** parameter specifies your system. More information for the volumes can be displayed as specified with the **-l** parameter.

```
tpctool> lsvol -user me -pwd mypass -url myhost:9550 -dev 2107.75DG000+0 -l
```

The following output is returned:

ID	Label	Size	Format
IBM. 2107-75DG000-111e+1+2107.75DG000+0	Sample_Vol_111E (ID:111e)	100	FB
IBM. 2107-75DG000-111d+1+2107.75DG000+0	Sample_Vol_111D (ID:111d)	100	FB
IBM. 2107-75DG000-111c+1+2107.75DG000+0	Sample_Vol_111C (ID:111c)	100	FB
FlashCopy	RealUsed	Encrypted	
UNAVAILABLE	100	No	
UNAVAILABLE	100	No	
UNAVAILABLE	100	No	
UDP1	UDP2	UDP3	
Sample	HostA	-	
Sample	HostA	-	
Sample	HostA	-	

FlashCopy

Indicates if the volume is in a FlashCopy relationship and whether it is a FlashCopy source or target. Volumes that are not in a FlashCopy relationship are displayed with a blank value. If information about the copy relationship of a volume is not available, UNAVAILABLE is displayed.

Note:

- This value is available for volumes of the following systems only: IBM® System Storage® DS8000® and IBM SAN Volume Controller.

Real Used

Indicates the amount of space, in gigabytes, that a volume uses. IBM Spectrum® Control allocates the entire space for standard-provisioned volumes when they are created. For thin-provisioned volumes, it does not. This column displays the space that is being used.

Note:

- If a system is new and there is no data in the volumes, the value is zero. For example, the Size field might show 16 GB but the Real Used is 0.
- For non-thin provisioned volumes, the Real Used value always matches the Size value.
- Thin provisioned volumes always have an asterisk before their name (Label).
- Thin provisioned volumes of other storage systems are not supported by IBM Spectrum Control. The Real Used value is displayed as N/A.

Encrypted

Indicates if the volumes are on encrypted disks (Yes or No).

UDP1

The user-defined property (UDP) contains anything of significance to the group administrator. This alphanumeric character string must contain fewer than 255 characters in length.

UDP2

The user-defined property (UDP) contains anything of significance to the group administrator. This alphanumeric character string must contain fewer than 255 characters in length.


UDP3

The user-defined property (UDP) contains anything of significance to the group administrator. This alphanumeric character string must contain fewer than 255 characters in length.

lsvolgroup

Use the **lsvolgroup** command to display a list of volume groups that are known to IBM Spectrum Control.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- svr server_id
(optional) The -svr parameter specifies that information is displayed for all volume groups that are associated with a specific server. You can use the **lssvr** command to return information, including the server IDs, for all servers that were discovered.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- l
Specifies that the long version of the information is listed, which includes the following information:
 - ID
 - Name
 - Server Name
 - Capacity (GB)
 - Used Capacity (GB)
 - Free Space (GB)

If this parameter is not issued, only the ID is returned.

Example: Displaying a list of volume groups

The following command generates a list of volume groups that are known to IBM Spectrum Control.

```
tpctool>lsvolgroup -svr c8bf94c4e7df11e1861cd2308c029302+hops3.storage.abc.ibm.com
+9.11.92.102+++ -l
```


The following output is returned:

ID	Name	Server Name	Capacity (GB)
Used Capacity (GB)	Free Space (GB)		
rootvg+hops3.storage.abc.ibm.com	rootvg	hops3.storage.abc.ibm.com	68.25
68.00	0.25		

lszone

Use the **lszone** command to list the zones in a zone set.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url

- Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric WWN
 - Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- active
 - Specifies that only the active zones are listed. If this option is not issued, all zones are listed.
- fs character
 - Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
 - Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
 - Lists help information for the command.
- qs character
 - Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
 - Suppresses all output for the command. If you omit this parameter, output is included.
- zone_set
 - Specifies the zone set. The *zone_set* variable is the name of the zone set.

Example: Listing all zones

The following command lists the names of all zones in the PARIS zone set:

```
tpctool> lszone -user me -pwd mypass -url myhost:myport
-fabric 100000051E34F6A8 PARIS
```

The following output is returned:

```
Name
=====
WINDOWSNT
SUNSOLARIS
TEST
...
```

Example: Listing only the active zones

The following command lists the active zones:

```
tpctool> lszone -user me -pwd mypass -url myhost:myport
-fabric 100000051E34F6A8 -active PARIS
```


The following output is returned:

```
Name
=====
SUNSOLARIS
```

lszs

Use the **lszs** command to list information about zone sets. This information includes the zone set name and status.

Syntax



Parameters and arguments

- user user_name
 - Specifies an IBM Spectrum Control user ID.
- pwd password
 - Specifies the password for the IBM Spectrum Control user ID.
- url url
 - Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric WWN
 - Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- active
 - Specifies that only information about the active zone set is listed. If you omit this option, information about all zone sets is listed.
- l
 - Specifies that the long version of the information is listed:
 - Name

- Status: active or inactive

If you omit this option, only the name of the zone is listed.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

Example: Listing the long version of information for all zone sets

The following command lists the long version of information for all zone sets:

```
tpctool> lszs -user me -pwd mypass -url myhost:myport
-fabric 100000051E34F6A8 -l
```

The following output is returned:

Name	Status
PARIS	ACTIVE
LONDON	INACTIVE

Example: Listing the active zone set

The following command lists the name of the active zone set:

```
tpctool> lszs -user me -pwd mypass -url myhost:myport
-fabric 100000051E34F6A8 -active
```

The following output is returned:

```
Name
=====
PARIS
```

mkappgroup

Use the **mkappgroup** command to create an application with a specific name, type, description, and user-defined properties. You must have Administrator authority to use this command.

Syntax

```
tpctool — mkappgroup — -user user_name -pwd password -url url — -name ApplicationGroupName →
      -description ApplicationDescription — -type type of application defined by the user →
      -subtype UserDefinedSubtype — -udp1 UDP1 — -udp2 UDP2 — -udp3 UDP3 →
```

Parameters and arguments

-user user_name

Specifies an IBM Spectrum Control user ID.

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-name ApplicationGroupName

A unique name not already in use, fewer than 60 characters long. It cannot contain any of the following characters:

```
\ / : * ? " < > | .
```

-description ApplicationGroupDescription

(optional) The description can contain any alphanumeric characters, and must be fewer than 255 characters long.

-type The type of application defined by the user

(optional) The type determines the type of application that is created. For example, a type might be Payroll, Finance, Project, General, or a similar type. This parameter can contain any alphanumeric character, and must be fewer than 128 characters in length.

-subtype UserDefinedSubtype

(optional) The subtype specifies another qualifier for the type. For example, if the -type is Db2® Group, the -subtype might be Logs, Tablespace, or another subtype that is related to Db2 Group. This parameter can contain any alphanumeric character, and must be fewer than 128 characters in length.

- udp1 UDP1
(optional) A user-defined property (UDP) that can contain anything of significance to the application administrator. The UDP can contain any alphanumeric character, and must be fewer than 255 characters in length.
- udp2 UDP2
(optional) A user-defined property that can contain anything of significance to the application administrator. The UDP can contain any alphanumeric character, and must be fewer than 255 characters in length.
- udp3 UDP3
(optional) A user-defined property that can contain anything of significance to the application administrator. The UDP can contain any alphanumeric character, and must be fewer than 255 characters in length.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Creating an ID, type and subtype for an application

The following command creates an application Accounting with a type of Finance and a subtype of Payroll:

```
tpctool> mkappgroup -name Accounting -type Finance -subtype Payroll
```

The following output is returned:

Name	ID	Status
Accounting	/Application/7506214	SUCCESS

Related reference

- modifyappgroup

mkdeptgroup

Use the **mkdeptgroup** command to create a department with a specific name, description, type, subtype, and user-defined properties. You must have Administrator authority to use this command.

Syntax

```
➤ tpcroot mkdeptgroup -useruser_name -pwdpassword -urlurl -name DepartmentGroupName ->
└─ -description DepartmentDescription -type The type of department defined by the user
└─ -subtype UserDefinedSubtype -udp1 UDP1 -udp2 UDP2 -udp3 UDP3
```

Parameters and arguments

- user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- name DepartmentGroupName
A unique name not already in use, with a maximum of 60 characters in length. It cannot contain any of the following characters:
`\ / : * ? " < > | .`
- description DepartmentGroupDescription
(optional) The description can contain any alphanumeric characters, with a maximum of 255 characters in length.
- type
(optional) The type of department that is created. This parameter can contain any alphanumeric character, with a maximum of 60 characters in length. For example, if a department is named, Accounting, then the type might be Finance.
- subtype UserDefinedSubtype
(optional) The subtype specifies another qualifier for the type. For example, if the -type is Finance, the -subtype parameter might be Payroll. This parameter can contain any alphanumeric character, with a maximum of 60 characters in length.
- udp1 UDP1

(optional) A user-defined property (UDP) that can contain anything of significance to the department administrator. The UDP can contain any alphanumeric character, with a maximum of 255 characters in length.

-udp2 UDP2

(optional) A user-defined property that can contain anything of significance to the department administrator. The UDP can contain any alphanumeric character, with a maximum of 255 characters in length.

-udp3 UDP3

(optional) A user-defined property that can contain anything of significance to the department administrator. The UDP can contain any alphanumeric character, with a maximum of 255 characters in length.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

Example: Creating an ID, type and subtype for a department

The following command creates the ID of /Department/5041, a type of Finance, with a subtype of Payroll for a department:

```
tpctool> mkdeptgroup -name Accounting -type Finance -subtype Payroll
```

The following output is returned:

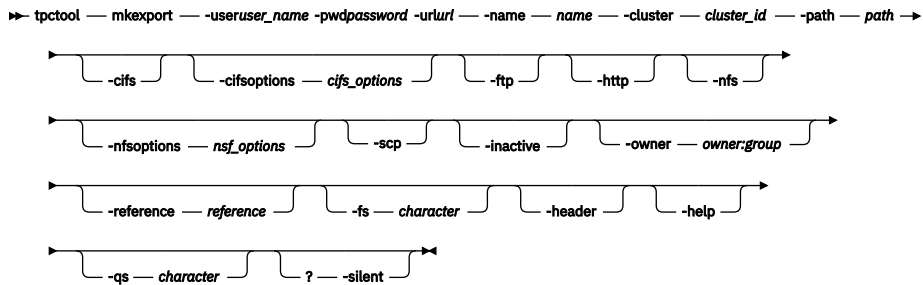
Name	ID	Status	Type	Subtype
Accounting	/Department/5041	Success	Finance	Payroll

mkexport

Use the **mkexport** command to create an export to access data through a data transfer protocol.

An *export* is a shared disk space that is accessible through the protocols that you specify when you run the **mkexport** command. You can create exports and enable them for HTTP, FTP, Secure Copy Protocol (SCP), Network File System (NFS), and Common Internet File System (CIFS) protocols. You must have Administrator authority to use this command.

Syntax



Parameters and arguments

-user user_name

Specifies an IBM Spectrum Control user ID.

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-name name

Specifies the name of the export.

-cluster cluster_id

Specifies the name of the IBM Spectrum® Control cluster key. The cluster key is listed in the ID column of the **lscluster** command output. This value is different from the Cluster ID output from the **lscluster** command.

-path path

Specifies the path for the export.

-cifs

Configures the CIFS protocol for the export.

-cifs options cifs_options

Defines the CIFS protocol options for the export. If the `cifsoptions` name or value contains spaces, the entire option must be enclosed in matching single quotation marks. The quotation marks must be preceded by an escape character.

- ftp
Configures FTP for the export.
- http
Configures HTTP for the export.
- nfs
Configures the NFS protocol for the export.
- nfsoptions `nfs_options`
Defines the NFS clients and their options for the export.
- scp
Configures SCP for the export.
- inactive
Marks the export as inactive. An inactive export is added to the list of exports, however you cannot access the data. You can use this parameter to modify the access control lists (ACLs) before you later activate the export with the **chexport** -active command.
- owner | -reference
Where owner `owner:group` sets the owner or owner group for the directory or reference `reference` sets the directory owner to the owner of the reference file. The reference file name must be an existing file or directory.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Creating an export that is enabled for HTTP, SCP, and NFS protocols

The following command creates an export to access data through HTTP, SCP, and the NFS protocol:

```
tpctool> mkexport -name cindyexport  
-cluster storage1.storage.tucson.ibm.com+storage1.storage.tucson.ibm.com+0  
-path /ibm/gpfs0/cindy -user admin -pwd password -url localhost:9550  
-http -scp -nfs
```

The following output is returned:

```
ExportId  
=====  
cindyexport+storage1.storage.tucson.ibm.com+storage1.storage.tucson.ibm.com+0  
  
Status  
=====  
SUCCESS
```

You are not required to enter credentials if you have already run the **tpctool** command with credentials.

Example: Creating an export that is enabled for CIFS protocol

The following command creates an export to access data through the CIFS protocol:

```
tpctool mkexport -name eexp10 -path /ibm/gpfs0/eexp10  
-cluster kq98n5d.ibm+00000200A22045DC+0 -cifs  
-cifsoptions "browseable=no,\"comment=comment for eexp10\",leases=no,  
sharemodes=no,syncio=yes,hideunreadable=yes,cifsacl=no,oplocks=no,  
locking=no,\"read only\",syncclose=no,\"access control=Everyone:ALLOWED:FULL;  
Administrator:ALLOWED:FULL\" -user db2admin -pwd g0vmw are -url localhost:9550
```

The following output is returned:

```
ExportId  
=====  
eexp10+kq98n5d.ibm+00000200A22045DC+0  
  
Status  
=====  
SUCCESS
```

Related reference

- [lscluster](#)
- [chexport](#)

Related information

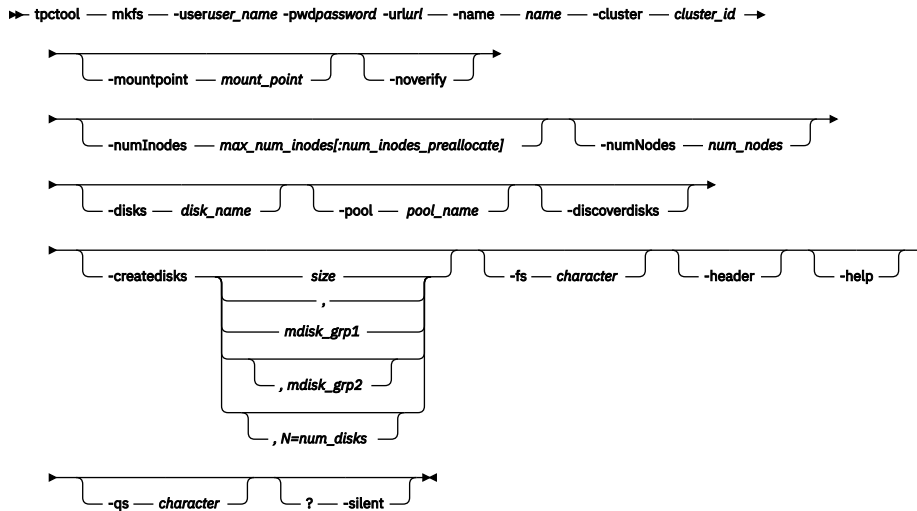
- http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp

mkfs

Use the **mkfs** command to create a GPFS file system to manage files on a storage device.

You can access the file system content by using file services like Common Internet File System (CIFS) or Network File System (NFS). You must have Administrator authority to use this command.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- name name
Specifies the name of the file system. The name must be unique and can be a maximum of 256 characters.
- cluster cluster_id
Specifies the name of the IBM Spectrum® Control cluster key. The cluster key is listed in the ID column of the **lscluster** command output.
Tip: This value is different from the `Cluster ID` output from the **lscluster** command.
- mountpoint mount_point
Specifies the mount point directory of the GPFS file system.
- noverify
Specifies that the disk descriptor is not verified so that disks that contain an old descriptor can be reused. Specify the `-noverify` parameter only when you want to reuse disks that are no longer needed for an existing file system. If the command is interrupted for any reason, you must use the `-noverify` parameter when you run the command again.
- numInodes max_num_inodes[:num_inodes_preallocate]
Specifies the maximum number of files for this file system. The *num_inodes_preallocate* variable specifies the number of inodes that the system immediately preallocates. You can specify values in thousands (k) or in millions (M). To specify values of 100 million for the *max_num_inodes* variable and 1million for the *num_inodes_preallocate* variable, enter `-numInodes 100M:1M`.
- numNodes num_nodes
Specifies the estimated number of nodes that is mounted with the file system.
- disks disk_name
Specifies the disks to create the file system on. The *disk_name* variable contains a comma-separated list of disk names.
Tip: You can verify the availability of a disk by running the **lsnsd** command.
- pool pool_name
Specifies the file system pool to create the file system on.
Tip: You can list all the disks in a pool by running the **lsnsd** command.
- discoverdisks
Specifies that IBM Spectrum Control detects and uses free GPFS Network Shared Disks (NSDs) automatically, which are tagged for a specified file system but not yet included. This option is applicable only for Storwize® V7000 Unified.
Tip: You can verify the availability of a disk by running the **lsnsd** command.
- createdisks size | mdisk_grp1 | mdisk_grp2 | N=num_disks
Creates disks implicitly, before the file system is created, and then adds them to the file system. This option is applicable only for Storwize V7000 Unified.
 - size
Specifies the size of the new disks. Size is specified as an integer with capacity up to petabyte without a space between the size and the unit; for example 17G. Disk sizes must be specified either without suffix (byte) or with K (kilobyte), M (megabyte), G (gigabyte), T (terabyte), or P (petabyte). Values less than 512 MB are not supported. This parameter is mandatory.
 - mdisk_grp1
Specifies the storage system managed disk (MDisk) group in which the underlying NAS volumes are created. This parameter is mandatory.

- Tip: You can see a list of available MDisk groups by using the **svcinfo lsmdiskgrp** command.
- mdisk_grp2**
Specifies the second storage system MDisk group in which the underlying NAS volumes are created. This parameter is optional.
- Tip: You can see a list of available MDisk groups by using the **svcinfo lsmdiskgrp** command.
- num_disks**
Specifies the number of storage system NAS volumes that is created in each MDisk group. This parameter is optional. The default number of disks is 3.
- fs character**
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header**
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?**
Lists help information for the command.
- qs character**
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent**
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Create a file system by using the createdisks parameter

The following command first creates one disk of 20 GB and then creates a file system.

```
tpctool> mkfs -name eefs05 -cluster kq98n5d.ibm+00000200A20045DC+0
-createdisks 20G,0,N=1 -user admin -pwd password -url localhost:9550
```

The following output is returned:

FilesystemId	Status
eefs05+kq98n5d.ibm+00000200A20045DC+0	SUCCESS

Example: Create a file system by using the pool parameter

The following command creates a file system on a specified pool.

```
tpctool> mkfs -name eefs06 -cluster kq98n5d.ibm+00000200A20045DC+0
-pool slpool -user admin -pwd password -url localhost:9550
```

The following output is returned:

FilesystemId	Status
eefs06+kq98n5d.ibm+00000200A20045DC+0	SUCCESS

Example: Create a file system by using the discoverdisks parameter

The following command discovers all the disks that are tagged to the **eefs07** file system, creates the **eefs07** file system, and finally adds the disks to it.

```
tpctool> mkfs -name eefs07 -cluster kq98n5d.ibm+00000200A20045DC+0
-discoverdisks -user admin -pwd password -url localhost:9550
```

The following output is returned:

FilesystemId	Status
eefs07+kq98n5d.ibm+00000200A20045DC+0	SUCCESS

Example: Create a file system by using the disks parameter

The following command creates a file system on the specified disks.

```
tpctool> mkfs -name eefs01 -cluster storage3.storage.tucson.ibm.com+9.11.92.174+0
-disks array0_sas_60001ff078c3a0789ff0001,array1_sas_60001ff078c3a0689fe0000
-user admin -pwd password -url localhost:9550
```

The following output is returned:

FilesystemId	Status
eefs01+storage3.storage.tucson.ibm.com+9.11.92.174+0	SUCCESS

Related reference

- [lscluster](#)
- [lsnsd](#)

Related information

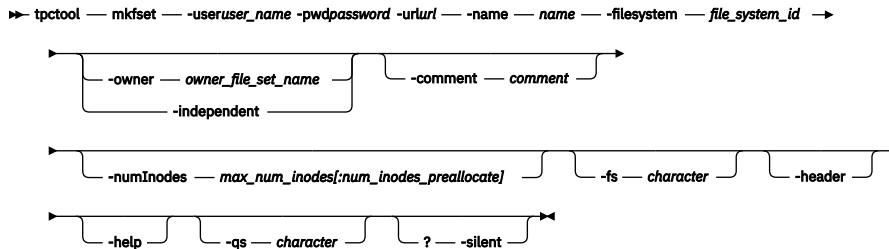
- http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp

mkfset

Use the **mkfset** command to create a fileset that is associated with a file system on a Storwize® V7000 File Module storage system. With filesets, you can use functions such as snapshots or quotas within a file system. You must have Administrator authority to use this command.

The new fileset is empty except for a root directory, and is not shown in the directory name space until you run the **linkfset** command. When you create a fileset, you can establish policies and quotas on the fileset before you link the fileset to the name space. You can create a maximum of 10,000 filesets for each file system.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- name name
Specifies the name of the fileset.
- filesystem file_system_id
Specifies the IBM Spectrum® Control file system key. The file system key is listed in the ID column of the **lsfs** command output.
- owner owner_file_set_name
Specifies the name of the independent fileset where the new file set is allocated. This parameter is mutually exclusive with the -independent parameter. Both filesets must belong to the same file system.
- independent
Specifies whether you want to create an independent fileset with its own allocated inodes. Otherwise, the fileset is created as a dependent fileset, and is allocated on the file system or in an independent fileset. This parameter is mutually exclusive with the -owner parameter.
- comment comment
Specifies a comment that displays in the output of the **lsfset** command. The length of this comment can be a maximum of 255 characters. You must enclose comments in double quotation marks.
- numInodes max_num_inodes[:num_inodes_preallocate]
Specifies the maximum number of files for a new, independent file set. This parameter can be used only with the -independent parameter. The *num_inodes_preallocate* variable specifies the number of inodes that the system immediately preallocates. You can specify values in thousands (k) or in millions (M). To specify values of 100 million for the *max_num_inodes* variable and 1million for the *num_inodes_preallocate* variable, enter `-numInodes 100M:1M`. GPFS defines a minimum number of inodes, which might be greater than the maximum specified. The default values for the fileset are 1million (1 M) for the *max_num_inodes* variable and 50,000 (50 K) for the *num_inodes_preallocate* variable.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Create a file system

The following command creates a fileset on a specified file system.

```
tpctool> mkfset -name eefset
-filesystem eefs+kq458mv.ibm+00000200A2A0153C+0
-user admin -pwd password -url localhost:9550
```

The following output is returned:

FilesetId	Status
eefset+eefs+kq458mv.ibm+00000200A2A0153C+0	SUCCESS

Related reference

- [linkset](#)
- [lsfs](#)
- [lsfset](#)

Related information

- <http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp>

mksrg

Use the **mksrg** command to create a storage resource group with a specific name, type, description, and user-defined properties. You must have Administrator authority to use this command.

Syntax

```
tpctool mksrg -user user_name -pwd password -url url -name SRGname -description SRGDescription
           -udp1 SRGUDP1 -udp2 SRGUDP2 -udp3 SRGUDP3 -fs character -header
           -help -qs character ? -silent
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- name SRGname
A unique name not already in use, fewer than 60 characters long. It cannot contain any of the following characters:
`\ / : * ? " < > | .`
The user name of the user who is creating the storage resource group is added to the beginning of the storage resource group name, followed by the name that is specified in the *name* variable.
- description SRGDescription
(optional) The description can contain any alphanumeric characters, and must be fewer than 255 characters long.
- udp1 SRGUDP1
(optional) A user-defined property that can contain anything of significance to the storage resource group administrator. The user-defined property (UDP) can contain any alphanumeric character, and must be fewer than 255 characters long.
- udp2 SRGUDP2
(optional) A user-defined property that can contain anything of significance to the storage resource group administrator. The UDP can contain any alphanumeric character, and must be fewer than 255 characters long.
- udp3 SRGUDP3
(optional) A user-defined property that can contain anything of significance to the storage resource group administrator. The UDP can contain any alphanumeric character, and must be fewer than 255 characters long.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Creating a storage resource group

The following command creates a storage resource group with the specified attributes:

```
tpctool> mksrg -name testsrg -description Test-SRG-Description
-udp1 myudp1 -udp2 myudp2 -udp3 myudp3
```

The following output is returned:

Name	Status
Administrator.testsrc	SUCCESS

mkzone

Use the **mkzone** command to create a zone. This command must be run within a transaction. You must have Administrator authority to use this command.

Important:

- An *orphan zone* is one that does not belong to a zone set. If the zone that you are creating is an orphan zone and the switch and data source that you are using support orphan zones, use only the **mkzone** command to create the zone. However, if the switch and data source that you are using do not support orphan zones, you must use the **addzone** command within the same transaction to add the zone to a zone set. Specify the **addzone** command AFTER you create the zone by using the **mkzone** command. For details about adding a zone to a zone set, see the **addzone** command.
- To activate your changes to a zone set, use the **actzs** command. For details, see the **actzs** command.

Syntax

➔ **tpctool** — **mkzone** — **-user** *user_name* **-pwd** *password* **-url** *url* — **-fabric** *WWN* — **-zone** *zone* — **-help**

└─ ? — **-silent** ─┐
└─ *ports* ─┐

Parameters and arguments

- user** *user_name*
Specifies an IBM Spectrum Control user ID.
- pwd** *password*
Specifies the password for the IBM Spectrum Control user ID.
- url** *url*
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric** *WWN*
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- zone** *zone*
Specifies the zone. The *zone* variable is the name of the zone.
- help** | **-h** | **-?**
Lists help information for the command.
- silent**
Suppresses all output for the command. If you omit this parameter, output is included.
- ports** | **-**
Specifies the switch ports. The *ports* variable is a list of worldwide port names (WWPNs). If you specify a single dash (-), the WWPNs are read from standard input.

Example: Creating a zone

In the following example, assume that the switch that you are using does not support orphan zones. To create the SUNSOLARIS zone, add the zone to a zone set, and activate the zone set. Specify all zoning changes within one transaction, from start to commit, and then activate the zone set by using the **actzs** command. The list of WWPNs is read from standard input.

```
tpctool> start -user me -pwd mypass -url myhost:myport -fabric 100000051E34F6A8
tpctool> mkzone -fabric 100000051E34F6A8 -zone SUNSOLARIS -
tpctool> addzone -fabric 100000051E34F6A8 -zs PARIS SUNSOLARIS
tpctool> commit -fabric 100000051E34F6A8
tpctool> actzs -fabric 100000051E34F6A8 PARIS
```

Related reference

- [addzone](#)
- [actzs](#)

mkzs

Use the **mkzs** command to create a zone set. This command must be run as a transaction. You must have Administrator authority to use this command.

Syntax

```

▶ tpctool — mkzs — -user user_name -pwd password -url url — -fabric WWN —
└─ -help ─┘ └─ ? — -silent ─┘

└─ zone_set ─▶

```

Parameters and arguments

-user user_name
Specifies an IBM Spectrum Control user ID.

-pwd password
Specifies the password for the IBM Spectrum Control user ID.

-url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).

-help | -h | -?
Lists help information for the command.

-silent
Suppresses all output for the command. If you omit this parameter, output is included.

zone_set
Specifies the zone set. The *zone_set* variable is the name of the zone set.

Tip: To create a zone set, you must include at least one zone.

Example: Creating a zone set

The following commands create the PARIS zone set, which includes the SUNSOLARIS zone:

```

tpctool> start -user me -pwd mypass -url myhost:myport -fabric 100000051E34F6A8
tpctool> mkzone -fabric 100000051E34F6A8 -zone SUNSOLARIS -
tpctool> mkzs -fabric 100000051E34F6A8 PARIS
tpctool> addzone -fabric 100000051E34F6A8 -zs PARIS SUNSOLARIS
tpctool> commit -fabric 100000051E34F6A8

```

modifyappgroup

Use the **modifyappgroup** command to modify an existing application. You can add and remove members of the application by specifying the member type and the key for the individual member. You can also use a -tagkey and -tagvalue member pair.

Syntax

```

▶ tpctool — modifyappgroup — -user user_name -pwd password -url url — -add — -remove — -id ApplicationGroupID —
└─ -type ─┘ └─ -key memberKey OR -tagkey ─┘ └─ -tagvalue tagvalue ─▶
├─ appgroup ─┘
├─ datastore ─┘
├─ export ─┘
├─ filesystem ─┘
├─ filesset ─┘
├─ server ─┘
├─ vm ─┘
├─ volume ─┘
└─ volumegroup ─┘
└─ udp1 ─┘
└─ udp2 ─┘
└─ udp3 ─┘

└─ -help ─┘ └─ -fs character ─┘ └─ -qs character ─┘ └─ ? — -silent ─┘ └─ -header ─▶

```

Parameters and arguments

-user user_name
Specifies an IBM Spectrum Control user ID.

-pwd password
Specifies the password for the IBM Spectrum Control user ID.

-url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-add
Adds a member to the specified application, which is designated by the specified member type and key or the -tagkey,-tagvalue parameter pair. Restriction: Vaults can only be added with the IBM Spectrum® Control GUI.

-remove
Removes a member from the specified application, which is designated by the member type and key that you specify or the -tagkey,-tagvalue parameter pair. Use the **lsappgroupmembers -id ApplicationGroupID** command to view a list of members and member keys for a specified application.

-id ApplicationGroupID

The identifier of the application that is being modified. Use the **lsappgroup** command, to view a list of the applications that are known to IBM Spectrum Control.

-type

appgroup | datastore | export | filesystem | fileset | server | vm |
volume | volumegroup

Denotes the type of element that needs to be added to or removed from the application.

-key memberKey

The unique key that corresponds to the member and member type that is specified by the **-type** parameter.
You can specify one of the following key types for this parameter:

appgroup

The identifier of the application to be added or removed.

datastore

The ID of the data store, which is composed of the data store name and the hypervisor host name.

To find the hypervisor server key for the data store, use the **lssvr** command, and the **lsfs** command, with the **-svr hypervisor_server_key** option, to obtain the ID of the data store. For example, C:\TPC_install_directory\cli>tpctool lssvr -l and C:\installation_dir\cli>tpctool lsfs -svr key_from_the_first output.

The following example shows how to add a data store to an application:

```
C:\installation_dir\cli>tpctool modifyappgroup -add -id  
ApplicationGroupID -type datastore -key ID_from_lsfs_output
```

export

The ID of the export, which is composed of the export name, cluster name, storage system name, and storage system format for IBM® Storwize® V7000 Unified and IBM Spectrum Scale. For servers and NetApp file servers, this parameter is composed of the export name and host name of the servers and NetApp file servers.

To find the ID for the export, use the **lssvr** command to locate the server key for servers and NetApp file servers, or use the **lsdev** command to get a share and the subsystem ID.

For example: C:\installation_dir\cli>tpctool lssvr -l and C:\installation_dir\cli>tpctool lsdev -subsys -l.

To find the cluster, run the following command:

```
C:\installation_dir\cli>tpctool lscluster -dev  
GUID_of_the_subsystem -l
```

The following example shows how to find the exports:

```
C:\installation_dir\cli>tpctool lsexport -svr  
server_key or C:\installation_dir\cli>tpctool  
lsexport -cluster ID_of_the_cluster
```

The following example shows how to add an export to an application:

```
C:\installation_dir\cli>tpctool modifyappgroup -add -id  
ApplicationGroupID -type export -key id_from_lsexport_output
```

Note: AIX® and Linux® exports are only available if the user of the computer explicitly exports the directories.

filesystem

The ID of the file system, which is composed of the file system name and host name of the server.

The ID of the file system is listed in the ID column of the **lsfs** command output.

To find the server key for the file system, use the **lssvr** command, and the **lsfs** command, with the **-svr server_key** option.

For example:

```
C:\installation_dir\cli>tpctool lssvr -l and C:\installation_dir\cli>tpctool lsfs -svr  
key_from_the_first_output
```

The following example shows how to add a file system to an application :

```
C:\installation_dir\cli>tpctool modifyappgroup -add -id  
ApplicationGroupID -type filesystem -key id_from_lsfs_command_output
```

fileset

The unique ID of the fileset, which is composed of the file system name, cluster name, storage system name, and storage system format for IBM Storwize V7000 Unified and IBM Spectrum Scale.

The unique ID of the fileset is listed in the File Set column of the **lsfset** command output.

The following example shows how to add a fileset to an application:

```
C:\installation_dir\cli>tpctool modifyappgroup -add -id  
ApplicationGroupID -type fileset -key id_from_lsfset_command_output
```

The following example shows how to remove a fileset from an application:

```
C:\installation_dir\cli>tpctool modifyappgroup -remove -id  
ApplicationGroupID -type fileset -key id_from_lsfset_command_output
```

server

The combination of the GUID, the host name, and the IP address. The unique key for the server is listed in the Key column of the **lssvr -l** command output.

The following example shows how to retrieve the ID for the server:

```
C:\installation_dir\cli>tpctool lssvr -l
```

The following example shows how to add a server to an application:

```
C:\installation_dir\cli>tpctool modifyappgroup -add -id
ApplicationGroupID -type server -key id_from_lssvr_output
```

vm

The ID of the virtual machine, which is composed of the virtual machine and the host name of the hypervisor. The ID of the virtual machine is listed in the Key column of the **lsvm** command output.

The following example shows how to obtain the ID for the virtual machine:

```
C:\installation_dir\cli>tpctool lsvm
```

The following example shows how to add a virtual machine to an application:

```
C:\installation_dir\cli>tpctool modifyappgroup -add -id
ApplicationGroupID -type vm -key id_from_lsvm_output
```

volume

The ID of the storage volume, which is composed of the storage volume name, storage system name, and storage system format. The ID of the storage volume is listed in the ID column of the **lsvol** command output.

The following example shows how to find the ID of the volume:

```
C:\installation_dir\cli>tpctool lsvol -dev
GUID_of_the_subsystem -l
```

The following example shows how to obtain the key for the subsystem of the volume:

```
C:\installation_dir\cli>tpctool lsdev -subsys -l
```

The following example shows how to add the ID of the volume to an application:

```
C:\installation_dir\cli>tpctool modifyappgroup -add -id
ApplicationGroupID -type volume -key volume_id
```

volumegroup

The ID of the volume group, which is composed of the volume group name and the host name. The ID of the volume group is listed in the Key column of the **lsvolgroup** command output.

The following example shows how to obtain the ID for the server:

```
C:\installation_dir\cli>tpctool lssvr -l
```

The following example shows how to find the ID of the volume group:

```
C:\installation_dir\cli>tpctool lsvolgroup -svr
server_ID_from_lssvr_output -l
```

The following example shows how to add the ID of the volume group to an application:

```
C:\installation_dir\cli>tpctool modifyappgroup -add -id
ApplicationGroupID -type volumegroup -key volumegroup_id
```

-tagkey udp1 or udp2 or udp3

Specifies the -tagkey parameter that is used to create application members.

Only the following -type parameters can be defined with tags:

- server
- filesystem
- appgroup

-tagvalue tagvalue

Specifies the -tagvalue parameter that is used to create application members.

Note: You can only specify one -tagkey and-tagvalue parameter member pair at a time.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

Example: Adding members to an application

The following example adds the application member that is specified by the **-type** and **-key** parameters, to /Application/5055.

You can add members to a specified application, with the following command:

```
tpctool
tpctool> modifyappgroup -add -id /Application/5044 -type appgroup
-key /Application/5055
```

The following output is returned:

Group Name	Member Key	Member Type	Status
DB2appgroup	/Application/5055	Application Group	SUCCESS

Example: Removing members of an application

The following example removes the application member that is specified by the **-type** and **-key** parameters, from application -id /Application/5055 .

You can remove members of a specified application with the following command:

```
tpctool
tpctool> modifyappgroup -remove -id /Application/5044 -type appgroup
-key /Application/5055
```

The following output is returned:

Group Name	Member Key	Member Type	Status
DB2appgroup	/Application/5055	Application Group	SUCCESS

Example: Adding application members with tags

The following example adds servers that are specified by the **-tagkey** and **-tagvalue** parameters to the application with -id /Application/5044.

You can add members to a specified application with tags using the following command:

```
tpctool
tpctool> modifyappgroup -id /Application/5044 -add -type server -tagkey
udp1 -tagvalue "dbservers"
```

The following output is returned:

Group Name	Member Key	Member Type	Status
DB2appgroup	bea358-ui.storage.abc.ibm.com+192.0.2.51+++	Server	SUCCESS
DB2appgroup	228e76-ui.storage.abc.ibm.com+192.0.2.76+++	Server	SUCCESS

Related reference

- [modifyappgroupviafile](#)

modifyappgroupviafile

Use the **modifyappgroupviafile** command to modify an existing application by using an input file. You can add and remove members by specifying the member type and identifiers or by using a member type with a -tagkey and -tagvalue pair.

You can specify -tagkey and -tagvalue pairs, for a given resource type, in the following ways:

- To select resources that match all of the specified tags, the -tagkey and -tagvalue parameters must be specified in this format:

```
-tagkey udp1 udp2 -tagvalue value1 value2
```

For example, to select the resources that have **udp1** set to **US** and **udp2** set to **West** specify:

```
-tagkey udp1 udp2 -tagvalue US West
```

- To select resources that match at least one of the specified tags, the -tagkey and -tagvalue parameters must be specified in this format:

```
-tagkey udp1 -tagvalue value1 -tagkey udp2 -tagvalue value2
```

For example, to select resources that have **udp1** set to **US** or **udp2** set to **West** specify:

```
-tagkey udp1 -tagvalue US -tagkey udp2 -tagvalue West
```

Syntax

```
➤ tpctool — modifyappgroupviafile — -user user_name -pwd password -url url — -inputfile — inputfilename ➤
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- inputfile inputfilename
Indicates the name of the input file, which contains information about the members that need to be added to or removed from applications.

All the other parameters are specified directly in the input file. The input file specifies the application where the members are added or removed by using the -add or -remove parameters:

```
-id <ApplicationGroupID> -add
-id <ApplicationGroupID> -remove
```


Then, the input file specifies the type and identifiers of the members for the application.
The -memberid parameter can be separated by spaces or commas. For example:

```
-memberid /Application/5044 /Application/5045  
-memberid /Application/5044,/Application/5045
```

If a -name parameter contains spaces or commas, it must be marked in quotations by using the double quotation character. For example:

```
-name "datastore1 (5) "
```

-id ApplicationGroupID

Indicates the unique ID of the application where the member is added or removed. Use the **lsappgroup** command, to view a list of the applications.

-type appgroup

Specifies that the entities are applications. The -memberid parameter specifies the unique IDs of the applications to be added or removed from the application, which is specified by the -id parameter. For example:

```
-type appgroup -memberid <appgroupname1>, <appgroupname2>, <appgroupname3>...
```

When you are using the -tagkey and -tagvalue member pair:

```
-type appgroup -tagkey <tag_key> -tagvalue <tag_value>
```

-type server

Specifies that the entities are servers or hypervisors. The -name parameter specifies the host names of the servers or hypervisors. For example:

```
-type server -name <hostname1>, <hostname2> ...
```

-type datastore

Specifies that the entities are data stores. The -servername parameter specifies the name of the hypervisor the data store belongs to and the -name parameter specifies the names of the data stores. For example:

```
-type datastore -servername <hostname> -name <datastore1> <datastore2> ...
```

-type export

Specifies that the entities are exports. The identifiers of the entities have different formats, which depend on whether the export members belong to servers or to IBM® Storwize® V7000 Unified cluster or IBM Spectrum Scale cluster. For example:

```
-type export -servername <hostname> -name <export1>, <export2> ..  
-type export -devicename <displaynameofdevice> -clustername <hostnameofcluster>  
-name <export1>, <export2>
```

-type filesystem

Specifies that the entities are file systems. The -servername parameter specifies the name of the server that the file system belongs to and the -name parameter specifies the names of the file system. For example:

```
-type filesystem -servername <hostname> -name <filesystem1>, <filesystem2> ..
```

When you are using the -tagkey and -tagvalue member pair:

```
-type filesystem -tagkey <tag_key> -tagvalue <tag_value>
```

-type fileset

Specifies that the entities are filesets of the file systems. The -devicename parameter specifies the IBM Storwize V7000 Unified or IBM Spectrum Scale storage system that the fileset belongs to. The -clustername parameter specifies the IBM Storwize V7000 Unified or IBM Spectrum Scale cluster the fileset belongs to. The -filesystemname parameter specifies the IBM Storwize V7000 Unified or IBM Spectrum Scale file system the fileset belongs to. For example:

```
-type fileset -devicename <displaynameofdevice>  
-clustername <hostnameofcluster> -filesystemname <filesystem1>  
-name <fileset1>, <fileset2>
```

-type vm

Specifies that the entities are virtual machines that were discovered by the hypervisor probe. The -servername parameter specifies the host name of the hypervisor that the virtual machine belongs to and the -name parameter specifies the names of the virtual machine. For example:

```
-type vm -servername <hostname> -name <vm1>, <vm2>, ..
```

-type volume

Specifies that the entities are storage volumes. The entities identifiers specify the display name of the storage system by using the -devicename parameter, and the display names of the storage volumes by using the -name parameter. For example:

```
-type volume -devicename <displaynameofdevice> -name <vol1>, <vol2>, ..
```

-type volumegroup

Specifies that the entities are the volume groups of the server. The entities identifiers specify the host name of the server and the names of the volume groups. For example:

```
-type volumegroup -servername <hostname> -name <vg1>, <vg2>, ...
```

-tagkey udp1 or udp2 or udp3

Specifies the tagkey that is used to create application members with an input file.
Only the following -type parameters can be defined with tags:

- server
- filesystem
- appgroup

-tagvalue tagvalue

Specifies the tagvalue that is used to create application members with an input file.

Example: Modifying members by using an input file

You can modify the members according to the contents of the input file with the following command:

```
tpctool
tpctool> modifyappgroupviafile -inputfile testModifyAppGroup04.txt
tpctool>
```

If the command successfully completes, no message is displayed.

If some members are not found in the IBM Spectrum® Control database, a warning message is displayed:

```
tpctool> modifyappgroupviafile -inputfile testModifyAppGroup05.txt
HWNDA0048W The following entities were not found: System1.storage.abc.ibm.com:
vm01,vm02,..
tpctool>
```

In this example, the `vm01` and `vm02` virtual machines of the `System1.storage.abc.ibm.com` hypervisor were not found in IBM Spectrum Control database.

Example: Adding and removing members from an application by using an input file

If you want to add and remove members from the application, `/Application/7599105`, specify two `-id` stanzas, one with the `-add` parameter and the other with the `-remove` parameter, followed by any other parameters. For example:

```
-id /Application/7599105 -add
    -type appgroup -memberid /Application/7605938
    -type export -servername server1 -name export1, export2, export3
    -type export -servername server2 -name export1, export2, export3
    -type export -servername server3 -name export1, export2, export3
    -type volume -servername server2 -name vg1, vg2
-id /Application/7599105 -remove
    -type appgroup -memberid /Application/7605939
    -type export -servername server1 -name export4, export5, export6
    -type volume -servername server2 -name vg3
```

Example: Adding and removing members from multiple applications by using an input file

To include more than one application, use multiple `-id` stanzas followed by other parameters. The same parameter can be used again for each `-id` stanza. For example:

```
-id /Application/7599105 -add
    -type appgroup -memberid /Application/7605938
    -type export -servername server1 -name export1, export2, export3
    -type volume -servername server2 -name vg1, vg2
-id /Application/7599105 -remove
    -type export -servername server1 -name export4, export5, export6
    -type volume -servername server2 -name vg3
-id /Application/7599106 -add
    -type vm -servername server1 -name vm1, vm2, vm3
    -type server -name server1, server2, server3
```

Example: Adding and removing members from an application by using an input file with tags

The following example adds application members that are specified by the `-tagkey` and `-tagvalue` parameter pairs. For example:

```
-id /Application/7599105 -add
    -type appgroup -tagkey udp1 -tagvalue TPCAppGroup
    -type server -tagkey udp1 -tagvalue Department1Servers
    -type filesystem -tagkey udp1 -tagvalue JFS
```

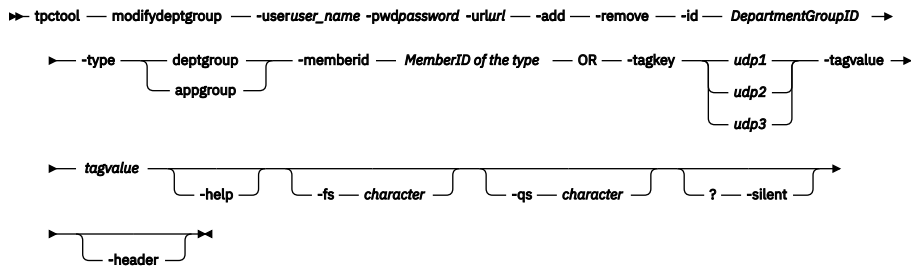
Related reference

- [modifyappgroup](#)

modifydeptgroup

Use the **modifydeptgroup** command to modify an existing department. You can add and remove members of a department by specifying the member type. For example, a deptgroup type for a department or an appgroup type for an application and the `-memberid` parameter for the individual member. You can also use a `-tagkey` and `-tagvalue` member pair.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
 - pwd password
Specifies the password for the IBM Spectrum Control user ID.
 - url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
 - add
Adds a member to the specified department, which is designated by the specified member -type and -memberid or the -tagkey, -tagvalue parameter pair.
 - remove
Removes a member from the specified department, which is designated by the specified member -type and -memberid parameters or the -tagkey, -tagvalue parameter pair. Use the **lsdeptgroupmembers -id id of the department** command to view a list of members and member IDs for a specified department.
 - id DepartmentGroupID
Indicates the unique ID of the department where the member is added or removed. Use the **lsdeptgroup** command, to view a list of the departments that are known to IBM Spectrum Control.
 - type deptgroup|appgroup
Denotes the department or application type that needs to be added or removed. This parameter can contain any alphanumeric character, with a maximum of 60 characters in length.
 - memberid MemberID of the type
The unique ID of the member being added or removed. Use the **lsdeptgroup** command, to view a list of the departments that are known to IBM Spectrum Control. You can specify one of the following key types for this parameter:
 - deptgroup
The unique name of the department to be added or removed that is specified by the **-memberid** parameter.
 - appgroup
The unique name of the application to be added or removed that is specified by the **-memberid** parameter.
 - tagkey udp1 or udp2 or udp3
Specifies the -tagkey parameter that is used to add or remove department members.
 - tagvalue tagvalue
Specifies the -tagvalue parameter that is used to add or remove department members.
- Note: You can only specify one -tagkey and -tagvalue parameter member pair at a time.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
 - header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
 - help | -h | -?
Lists help information for the command.
 - qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
 - silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Adding members to a department

The following example adds the department member that is specified by the **-type** and **-memberid** parameters, with /Department/5044 -type deptgroup.

You can add members of a specified department, with the following command:

```

tpctool
tpctool> modifydeptgroup -id /Department/5041 -add -type deptgroup
-memberid /Department/5044

```

The following output is returned:

Group Name	Member ID	MemberType	Status
Accounting	/Department/5044	Department	SUCCESS

You can add an appgroup type to a specified department, with the following command:

```

tpctool
tpctool> modifydeptgroup -id /Department/5035 -add -type appgroup
-memberid /Application/7024

```

The following output is returned:

Group Name	Member ID	Member Type	Status
Sales	/Application/7024	Application	SUCCESS

Example: Removing members of a department

The following example removes the department member that is specified by the **-type** and **-memberid** parameters.

You can remove members of a specified department with the following command:

```
tpctool
tpctool> modifydeptgroup -id /Department/5041 -remove -type deptgroup
-memberid /Department/5044
```

The following output is returned:

Group Name	Member ID	MemberType	Status
Accounting	/Department/5044	Department	SUCCESS

You can remove the appgroup type for a specified department, with the following command:

```
tpctool
tpctool> modifydeptgroup -id /Department/5035 -remove -type appgroup
-memberid /Application/7024
```

The following output is returned:

Group Name	Member ID	Member Type	Status
Sales	/Application/5044	Application	SUCCESS

Example: Adding department members with tags

The following example adds department members that are specified by the **-tagkey** and **-tagvalue** parameters.

You can add members of a specified department with tags using the following command:

```
tpctool
tpctool> modifydeptgroup -add -id /Department/5041 -type deptgroup -tagkey udp1
-tagvalue value1
```

The following output is returned:

Group Name	Member ID	Member Type	Status
Accounting	/Department/5044	Department Group	SUCCESS
Accounting	/Department/5035	Department Group	SUCCESS
Accounting	/Department/6044	Department Group	SUCCESS
Accounting	/Department/6035	Department Group	SUCCESS

Example: Removing department members with tags

The following example removes department members that are specified by the **-tagkey** and **-tagvalue** parameters.

You can remove members of a specified department with tags using the following command:

```
tpctool
tpctool> modifydeptgroup -remove -id /Department/5041 -type deptgroup -tagkey udp1
-tagvalue value1
```

The following output is returned:

Group Name	Member ID	Member Type	Status
Accounting	/Department/5044	Department Group	SUCCESS
Accounting	/Department/5035	Department Group	SUCCESS
Accounting	/Department/6044	Department Group	SUCCESS
Accounting	/Department/6035	Department Group	SUCCESS

Related reference

- [modifydeptgroupviafile](#)

modifydeptgroupviafile

Use the **modifydeptgroupviafile** command to modify an existing department by using an input file. You can add and remove members by specifying the member type (deptgroup, appgroup) and identifiers or by using a member type with a **-tagkey** and **-tagvalue** pair.

You can specify **-tagkey** and **-tagvalue** pairs, for a given resource type, in the following ways:

- To select resources that match all of the specified tags, the **-tagkey** and **-tagvalue** parameters must be specified in this format:

```
-tagkey udp1 udp2 -tagvalue value1 value2
```

For example, to select the resources that have `udp1` set to `US` and `udp2` set to `West` specify:

```
-tagkey udp1 udp2 -tagvalue US West
```

- To select resources that match at least one of the specified tags, the `-tagkey` and `-tagvalue` parameters must be specified in this format:

```
-tagkey udp1 -tagvalue value1 -tagkey udp2 -tagvalue value2
```

For example, to select resources that have `udp1` set to `US` or `udp2` set to `West` specify:

```
-tagkey udp1 -tagvalue US -tagkey udp2 -tagvalue West
```

Syntax

```
➔ tpctool — modifydeptgroupviafile — -user user_name -pwd password -url url — -inputfile — inputfilename ➔
```

Parameters and arguments

`-user user_name`

Specifies an IBM Spectrum Control user ID.

`-pwd password`

Specifies the password for the IBM Spectrum Control user ID.

`-url url`

Specifies the Device server. The format of the URL is `system:port_number`, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

`-inputfile inputfilename`

Indicates the name of the input file that is used to parse the input file from the CLI program.

All the other parameters are specified directly in the input file. The input file specifies the department where the members are added or removed by using the `-add` or `-remove` parameters:

```
-id <deptgroupname> -add
```

```
-id <deptgroupname> -remove
```

Then, the input file specifies the type and identifiers of the members for the department.

The `-memberid` parameter can be separated by spaces or commas. For example:

```
-memberid /Department/5044 /Department/5045
```

```
-memberid /Department/5044,/Department/5045
```

`-id DepartmentGroupID`

Indicates the unique ID of the department where the member is added or removed. Use the `lsdeptgroup` command, to view a list of the departments that are known to IBM Spectrum Control.

`-type deptgroup|appgroup`

Specifies that the entities are a department or application. The `-memberid` parameter specifies the unique names of the departments or applications to be added or removed from the department, which is specified by the `-id` parameter. For example:

```
-type deptgroup -memberid /Department/5045, /Department/5046,  
/Department/5047....
```

```
-type appgroup -memberid /Application/5050, /Application/5051,  
/Application/5052....
```

When you are using the `-tagkey` and `-tagvalue` member pair:

```
-type deptgroup -tagkey <tag_key> -tagvalue <tag_value>
```

`-memberid MemberID of the type`

The unique ID of the member being added or removed. Use the `lsdeptgroup` command, to view a list of the departments that are known to IBM Spectrum Control. Use the `lsappgroup` command to view the list of applications.

You can specify one of the following types of ids for this parameter:

- Department -memberid, for example, `/Department/12345`.
- Application -memberid, for example, `/Application/56789`.

`-tagkey udp1 or udp2 or udp3`

Specifies the `-tagkey` parameter that is used to add or remove department members, with an input file.

`-tagvalue tagvalue`

Specifies the `-tagvalue` parameter that is used to add or remove department members, with an input file.

When you are using the `-tagkey` and `-tagvalue` member pair:

```
-type deptgroup -tagkey <tag_key> -tagvalue <tag_value>
```

Example: Modifying department members by using an input file

You can modify the department members according to the contents of the input file with the following command:

```
tpctool  
tpctool> modifydeptgroupviafile -inputfile testModifyDeptGroup04.txt  
tpctool>
```

If the command successfully completes, no message is displayed.

If some members are not found in the IBM Spectrum Control database, a warning message is displayed:

```
tpctool> modifydeptgroupviafile -inputfile testModifyDeptGroup04.txt
HWNDA0048W The following entities were not found: /Department/5041,
/Department/5044, /Department/7044,..
tpctool>
```

Example: Adding and removing members from a department by using an input file

If you want to add and remove members from the deptgroup1 department, specify two `-deptgroupname` stanzas, one with the `-add` parameter and the other with the `-remove` parameter, followed by any other parameters. For example:

```
-id /Department/5041 -add
-type deptgroup -memberid /Department/6000, /Department/6001, /Department/6002
-type appgroup -memberid /Application/7000, /Application/7001, /Application/7002
-remove
-type deptgroup -memberid /Department/6000, /Department/6001, /Department/6002
```

Example: Adding and removing members from multiple departments by using an input file

To include more than one department, use multiple `-id` stanzas followed by other parameters. The same parameter can be used again for each `-id` stanza. For example:

```
-id /Department/5041 -add
-type appgroup -memberid /Application/7004, /Application/7005, /Application/7006,
/Application/7007, /Application/7008, /Application/7009

-id /Department/5044 -remove
-type appgroup -memberid /Application/7004, /Application/7005, /Application/7006,
/Application/7007, /Application/7008, /Application/7009

-id /Department/7044 -add
-type deptgroup -memberid /Department/6000, /Department/6001, /Department/6002,
/Department/6003, /Department/6004, /Department/6005
```

Example: Adding and removing members from an department by using an input file with tags

The following example adds department members that are specified by the `-tagkey` and `-tagvalue` parameter pairs. For example:

```
-id /Department/9003 -add
-type deptgroup -tagkey udp1 -tagvalue TPCDeptGroup
-type appgroup -tagkey udp1 -tagvalue TPCAppGroup
```


Related reference

- [modifydeptgroup](#)

modifysrg

Use the **modifysrg** command to modify an existing storage resource group. You can add and remove members by specifying the member type, such as switch or volume, and the key for the member you want to add or remove.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- add
Adds a member to the specified storage resource group, which is designated by the specified member type and key.
- remove
Removes an existing member from the specified storage resource group, which is designated by the specified member type and key. Use the **lssrgmembers -name SRGName** to see a list of members and member keys for a specified storage resource group.
- name SRGName
Indicates the unique storage resource group where the member is added or removed. To view a list of the storage resource groups that are known to IBM Spectrum Control, see the output of the **lssrg**.
- type
srg | fabric | switch | server | subsystem | pool |
volume | hostdisk | vm disk | mdisk | disk | filesystem | export | vm | volume group | datastore
Denotes the type of element that corresponds to the unique key in the **key** parameter of this function.
- key memberKey
Specifies the unique key that corresponds to the member and member type that is specified in the **type** parameter.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Adding or removing members of a storage resource group

You can add or remove members of a specified storage resource group with the following command:

```
tpctool
tpctool> modifysg -remove -name test.srg -type srg
-key Administrator.my-new-example-srg
```


The following output is returned:

SRG Name	Member Key	Member Type	Status
test.srg	Administrator.my-new-example-srg	Storage Resource Group	SUCCESS

mountfs

Use the **mountfs** command to mount a file system on all interface and management nodes or a specified subset. You must have Data Administrator authority to use this command.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- filesystem file_system_id
Specifies the ID of file system to be mounted.
- nodes nodes

Lists the nodes to mount the file system on in a comma-separated list. Specify only the interface node, management node, or both. If you omit this parameter, the file system is mounted on all interface and management nodes.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

Example: Mount a file system

The following command mounts a file system on a management node.

```
tpctool> mountfs -filesystem eefs+kq458mv.ibm+00000200A2A0153C+0  
-user admin -pwd password -url localhost:9550 -nodes mgmt001st001
```

The following output is returned:

FilesystemId	Status
eefs+kq458mv.ibm+00000200A2A0153C+0	SUCCESS

Related information

- ➦ http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp

rmappgroup

Use the **rmappgroup** command to delete a specific application and its members. If you use the **rmappgroup** command without the **-rmchildren** parameter, the application is deleted, and the members are moved up a level in the hierarchy.

Syntax

```
➦ tpctool — rmappgroup — -user user_name -pwd password -url url — -id ApplicationGroupID — -rmchildren —  
-fs character -header -help -qs character ? -silent
```

Parameters and arguments

-user user_name

Specifies an IBM Spectrum Control user ID.

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-id ApplicationGroupID

Specifies the unique ID of the application to be deleted. For example, **/Application/5044** where **5044** is the database ID of the application.

-rmchildren

(Optional) Specifies that any members of the application can be deleted. When you use this parameter, any subcomponent that belongs only to the application is removed. If the subcomponent belongs to more than one application, it is removed from the application and not deleted.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

Example: Deleting an application and the subcomponents

The following command deletes a specified application and the subcomponents:

```
tpctool
tpctool> rmappgroup -id /Application/5044 -rmchildren
```

The following output is returned:

Name	Status
Db2appgroup	SUCCESS

Example: Deleting an application and moving the subcomponent up a level

The following command deletes a specified application and moves the subcomponent up a level in the hierarchy:

```
tpctool
tpctool> rmappgroup -id /Application/5044
```

The following output is returned:

Name	Status
Db2appgroup	SUCCESS

If there are no top level applications in the hierarchy, the subcomponents become the top level application.

rmbackenddisktype

Use the **rmbackenddisktype** command to remove a back-end type of disk.

This command is available for the following storage systems:

- Storwize® V7000
- Storwize V7000 Unified
- SAN Volume Controller

You must have Administrator authority to use this command.

The following actions occur when you issue the **rmbackenddisktype** command to remove a back-end type of disk.

- The type is removed from the set of back-end types of disk.
- The type is reset to *DEFAULT* in back-end storage subsystems that are set to the type of disk that is removed.

Tip: You can use the following commands to set the back-end type of disk:

- **setarray**
- **setbackenddisktype**

You can also set the back-end type of disk for storage systems on the MDisk Group Details page in the IBM Spectrum® Control graphical user interface.

Syntax

```
tpctool rmbackenddisktype -type disk_type -user user_name -pwd password -url url -fs character -header -help -qs character ? -silent -grouping
```

Parameters and arguments

-type disk_type

Specifies the back-end type of disk. Specify a back-end type of disk or enter one of the following values.

Value	Description
A07	Sata - 7,500 rpm
F10	Fiber - 10,000 rpm
F15	Fiber - 15,000 rpm
DEFAULT	Default value

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

-silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Removing a back-end type of disk


```
tpctool> rmbackenddisktype -type A07
```

```
Back-end Disk Type Status
=====
A07                      Succeeded
```

rmbackendraidtype

- Storwize® V7000
- SAN Volume Controller

Syntax



Parameters and arguments

Value	Description
DEFAULT	The default RAID type.
1	RAID 1
5	RAID 5
6	RAID 6
X	The RAID type that is used by IBM® XIV® Storage System.

Example: Removing a type of back-end RAID array

The following command removes a type of back-end RAID array:

```
tpctool> rmbackendraidtype -type testing
```

If the command is successful, a list of the types of back-end RAID types that are removed is displayed:

```
Back-end Type Status
=====
testing          Succeeded
```

rmbackendtype

Use the **rmbackendtype** command to remove a type of back-end storage system.

This command is available for the following storage systems:

- Storwize® V7000
- Storwize V7000 Unified
- SAN Volume Controller

You must have Administrator authority to use this command.

When you issue the **rmbackendtype** command to remove a type of back-end storage system, the following actions occur:

- The type is removed from the set of back-end types.
- The type is reset to *DEFAULT* in back-end storage systems that are set to the back-end type that is removed.

Tip: You can use the following commands to set the type of back-end storage system:

- **setarray**
- **setbackendtype**

You can also set the type of back-end storage system on the MDisk Group Details page in the IBM Spectrum® Control graphical user interface.

Syntax

```
➤ tpctool — rmbackendtype — -type — storage_system_type — -user user_name -pwd password -url url — -fs — character —
-headers — -help — -qs — character — ? — -silent — -grouping —
```

Parameters and arguments

-type storage_system_type

Specifies the type of back-end storage system that manages most of the storage pool resources. Specify a type of back-end storage system or enter one of the following values.

Value	Name
C	EMC Clariion
D	IBM® System Storage® DS8000®
S	EMC Symmetrix
X	IBM XIV® Storage System
DEFAULT	Default value

Tip: Use the **lsbackendtypes** command to show a list of back-end storage systems.

-user user_name

Specifies an IBM Spectrum Control user ID.

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

-grouping

Enables the grouping of numeric values. For example, in English the value 12000 would display as 12,000. The grouping character is determined by the system locale.

Example: Removing a type of back-end storage system

The following command removes the specified type of back-end storage system:

```
tpctool> rmbackendtype -type L
```

The type of back-end storage system that is removed is displayed:

```
Back-end Type Status
=====
L              Succeeded
```

rmdeptgroup

Use the **rmdeptgroup** command to delete a specific department and its members. If you use the **rmdeptgroup** command without the **-rmchildren** parameter, the department is deleted, and the members are moved up a level in the hierarchy.

Syntax

➔ **tpctool** — **rmdeptgroup** — **-user** *user_name* **-pwd** *password* **-url** *url* — **-id** *DepartmentGroupID* — **-rmchildren**

└─ **-fs** *character* ─┘ └─ **-header** ─┘ └─ **-help** ─┘ └─ **-qs** *character* ─┘ └─ **?** **-silent** ─┘

Parameters and arguments

- user** *user_name*
Specifies an IBM Spectrum Control user ID.
- pwd** *password*
Specifies the password for the IBM Spectrum Control user ID.
- url** *url*
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- id** *DepartmentGroupID*
Specifies the ID of the department to be deleted.
- rmchildren**
(Optional) Specifies that any members of the department are be deleted. When you use this parameter, any member that belongs only to the department is removed. If the members belong to more than one department, they are removed from the department and not deleted.
- fs** *character*
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header**
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help** | **-h** | **-?**
Lists help information for the command.
- qs** *character*
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent**
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Deleting a department and the subdepartments

The following command deletes a specified department and the subdepartments:

```
tpctool
tpctool> rmdeptgroup -id /Department/5041 -rmchildren
```

The following output is returned:

```
Name              Status
=====
Accounting        SUCCESS
```

Example: Deleting a department and moving the subdepartment up a level

The following command deletes a specified department and moves any subdepartments up one level in the hierarchy:

```
tpctool
tpctool> rmdeptgroup -id /Department/5041
```

The following output is returned:

```
Name              Status
=====
Accounting        SUCCESS
```

If there are no top level departments in the hierarchy, the subdepartment or subdepartments become the top level department.

Related reference

- [mkdeptgroup](#)

rmexport

Use the **rmexport** command to remove an export.

Removing an export does not remove the data or folder from the file system. You must have Administrator authority to use this command.

When you remove an export, the system behaves differently according to the export protocols. When an NFS export is removed, all connections to that shared space are immediately closed. When a CIFS export is removed, any user who is connected to the shared space can continue to write to that space. However, if the user disconnects, reconnects, and then attempts to write to the space again, the write operation fails.

Syntax

```
tpctool --rmexport --user user_name -pwd password -url url --export export_id -fs character
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- export export_id
Specifies the name of the export. This is the export key that is listed in the ID column of the **lsexport** command output.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Removing an export

The following command deletes the specified export:

```
tpctool> rmexport -export cindyexport+storage1.storage.tucson.ibm.com+storage1.storage.tucson.ibm.com+0 -user admin -pwd password -url localhost:9550
```

The following output is returned:

```
ExportId
=====
cindyexport+storage1.storage.tucson.ibm.com+storage1.storage.tucson.ibm.com+0

Status
=====
SUCCESS
```

Related reference

- [lsexport](#)

Related information

- http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp


rmfs

Use the **rmfs** command to remove a file system from an active management node. You must have Administrator authority to use this command.

Attention: When you remove a file system, all the data on that file system is deleted.

You must unmount a file system on all nodes before you can delete it. Use the **unmountfs** command to unmount a file system.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- filesystem file_system_id
Specifies the ID of the file system to be removed.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Remove a file system

The following command removes a file system.

```
tpctool> rmfs -filesystem eefs+kq458mv.ibm+00000200A2A0153C+0  
-user admin -pwd password -url localhost:9550
```

The following output is returned:

FilesystemId	Status
eefs+kq458mv.ibm+00000200A2A0153C+0	SUCCESS

Related reference

- [unmountfs](#)

Related information


- http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp

rmfset

Use the **rmfset** command to remove a fileset from a file system. You must have Administrator authority to use this command.

Important: You must use the **unlinkfset** command to unlink a fileset from a file system before you can remove it.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password

- Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- help | -h | -?
Lists help information for the command.
- fileset file_set_id
Specifies the IBM Spectrum® Control key of the fileset to be removed. The fileset key is listed in the **ID** column of the **lsfset** command output.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

Example: Removing a fileset

The following command removes the fileset named **eefset01+eefs+kq458mv.ibm+00000200A2A0153C+0**.

```
tpctool> rmfset -fileset eefset01+eefs+kq458mv.ibm+00000200A2A0153C+0
-user admin -pwd password -url localhost:9550
```

The following output is returned:

FilesetId	Status
eefset01+eefs+kq458mv.ibm+00000200A2A0153C+0	SUCCESS

Related reference

- [unlinkfset](#)
- [lsfset](#)

Related information

- http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp

rmsrg

Use the **rmsrg** command to delete the specified Storage Resource Group. The group types (Storage Resource Groups, Reporting Groups and Application Groups) that have child groups are not deleted. A warning message is issued and indicates that the group type being deleted contains child groups and the child groups must be deleted before the parent group can be deleted.

Syntax

```
tpctool — rmsrg — -user user_name -pwd password -url url — -name SRGname — -fs character —
-headers — -help — -qs character — ? — -silent
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- name SRGname
Specifies the name of the Storage Resource Group to be deleted.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

Example: Deleting a Storage Resource Group

The following command deletes a specified Storage Resource Group:

```
tpctool
tpctool> rmsrg -name Administrator.testsg
```

The following output is returned:

Name	Status
=====	
Administrator.testsg	SUCCESS

rmza

Use the **rmza** command to remove a zone alias or aliases from a zone. You must have Administrator authority to use this command.

Syntax

```
tpctool — rmza — -user user_name -pwd password -url url — -fabric WWN — -zone zone —
                                     -help — ? — -silent —
zone — zone_alias —
```

Parameters and arguments

-user user_name
Specifies an IBM Spectrum Control user ID.

-pwd password
Specifies the password for the IBM Spectrum Control user ID.

-url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).

-help | -h | -?
Lists help information for the command.

-silent
Suppresses all output for the command. If you omit this parameter, output is included.

-zone zone
Specifies the zone. The *zone* variable is the name of the zone.

zone_alias
Specifies the name or names of the zone aliases to be removed from the zone.

Example: Removing zone aliases from a zone

The following commands remove the PARIS zone alias from the EUROPE zone:

```
tpctool -user me -pwd mypass -url myhost:myport
tpctool> start -fabric 100000051E34F6A8
tpctool> rmza -fabric 100000051E34F6A8 -zone EUROPE PARIS
tpctool> commit -fabric 100000051E34F6A8
```

rmzaports

Use the **rmzaports** command to remove a port or ports from a zone alias. You must have Administrator authority to use this command.

Syntax

```
tpctool — rmzaports — -user user_name -pwd password -url url — -fabric WWN —
                                     -help — ? — -silent —
-za — zone_alias — port —
```

Parameters and arguments

-user user_name

- id *id*
Specifies an IBM Spectrum Control user ID.
- pwd *password*
Specifies the password for the IBM Spectrum Control user ID.
- url *url*
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric *WWN*
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- help | -h | -?
Lists help information for the command.
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- za *zone_alias*
Specifies the name of the zone alias from which the ports to be removed.
- port*
Specifies the name or names of the ports to be removed from the zone alias.

Example: Removing a port from a zone alias


The following commands remove a port from the PARIS zone alias:

```
tpctool> -user me -pwd mypass -url myhost:myport
tpctool> start -fabric 100000051E34F6A8
tpctool> rmzapsorts -fabric 100000051E34F6A8 -za PARIS 210000E08B0B4C2G
tpctool> commit -fabric 100000051E34F6A8
```

rmzone

Use the **rmzone** command to delete a zone or remove a zone from a zone set. If you remove or delete the last zone in a zone set, the zone set is also deleted. This command must be run as a transaction. You must have Administrator authority to use this command.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- zone zone
Specifies the zone. The *zone* variable is the name of the zone. If the last zone in the zone set is removed or deleted, the zone set also is deleted.
- zs zone_set
Specifies that zone is removed from the zone set. The *zone_set* variable is the name of the zone set. If this option is not issued, the zone is deleted.
- help | -h | -?
Lists help information for the command.
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Tip: To delete a zone from a zone set, the zone must be included in another zone set.

Example: Deleting a zone

The following commands delete the `WINDOWSNT` zone:

```
tpctool> start -user me -pwd mypass -url myhost:myport -fabric 100000051E34F6A8
tpctool> rmzone -fabric 100000051E34F6A8 -zone WINDOWSNT
tpctool> commit -fabric 100000051E34F6A8
```

Example: Removing a zone from a zone set

The following commands remove the `WINDOWSNT` zone from the `PARIS` zone set:

```
tpctool> start -user me -pwd mypass -url myhost:myport -fabric 100000051E34F6A8
tpctool> rmzone -fabric 100000051E34F6A8 -zone WINDOWSNT -zs PARIS
tpctool> commit -fabric 100000051E34F6A8
```

rmzoneports

Use the **rmzoneports** command to remove switch ports from a zone. This command must be run as a transaction. You must have Administrator authority to use this command.

Syntax

► **tpctool** — **rmzoneports** — **-user***user_name* **-pwd***password* **-url***url* — **-fabric** — *WWN* — **-zone** — *zone* — **-help** —

— **?** — **-silent** — **ports** — ►

Parameters and arguments

-user *user_name*
Specifies an IBM Spectrum Control user ID.

-pwd *password*
Specifies the password for the IBM Spectrum Control user ID.

-url *url*
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fabric *WWN*
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).

-zone *zone*
Specifies the zone. The *zone* variable is the name of the zone.

-help | **-h** | **-?**
Lists help information for the command.

-silent
Suppresses all output for the command. If you omit this parameter, output is included.

ports | **-**
Specifies the switch ports. The *ports* variable is a list of worldwide port names (WWPNs). If you specify a single dash (-), the WWPNs are read from standard input.

Example: Removing switch ports from a zone set

The following commands remove several switch ports from the SUNSOLARIS zone. The list of WWPNs is read from standard input:

```
tpctool> start -user me -pwd mypass -url myhost:myport -fabric 100000051E34F6A8
tpctool> rmzoneports -fabric 100000051E34F6A8 -zone SUNSOLARIS -
tpctool> commit -fabric 100000051E34F6A8
```

rmzs

Use the **rmzs** command to delete a zone set. This command must be run as a transaction. You must have Administrator authority to use this command.

Syntax

► **tpctool** — **rmzs** — **-user***user_name* **-pwd***password* **-url***url* — **-fabric** — *WWN* — **-help** — **?** — **-silent** — *zone_set* — ►

Parameters and arguments

-user *user_name*
Specifies an IBM Spectrum Control user ID.

-pwd *password*
Specifies the password for the IBM Spectrum Control user ID.

-url *url*
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fabric *WWN*
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).

-help | **-h** | **-?**
Lists help information for the command.

-silent
Suppresses all output for the command. If you omit this parameter, output is included.

zone_set
Specifies the zone set. The *zone_set* variable is the name of the zone set.

Tip: To delete a zone set, it must be inactive. If you attempt to delete a zone set that is empty, the zone set is not deleted and no error message is given.

Example: Deleting a zone set

The following commands delete the PARIS zone set:

```
tpctool> start -user me -pwd mypass -url myhost:myport -fabric 100000051E34F6A8
tpctool> rmzs -fabric 100000051E34F6A8 PARIS
tpctool> commit -fabric 100000051E34F6A8
```

rollback

Use the **rollback** command to erase any commands that were issued since you started the transaction. You must have Administrator authority to use this command.

Syntax

```
➤ tpctool — rollback — -user user_name -pwd password -url url — -fabric — WWN — -help —
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric WWN
Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- help | -h | -?
Lists help information for the command.

Example: Rolling back a transaction

The following command rolls back a transaction. The user has previously issued connection options, started a transaction, and issued a fabric-control command:

```
tpctool> rollback -fabric 100000051E34F6A8
```

runoptschedule

Use the **runoptschedule** command to run a schedule that analyzes storage tiering.

In the GUI, you can create schedules that are based on the criteria that you entered for tiering analysis. Each time that the schedule is run, a job with a unique ID is created.

When you issue the **runoptschedule** command, the job ID for the schedule is shown.

Tip: Issue the **lsopschedules** command to show a list of schedules and schedule IDs. After you issue the **runoptschedule** command, you can issue the **showoptresults** command to show the results of the analysis.

Syntax

```
➤ runoptschedule — -schedule_id schedule_id — -user user_name -pwd password -url url — -fs — character —  
-header -help -qs — character
```

Parameters and arguments

- schedule_id schedule_id
Specifies the ID of the schedule.
- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").

Example: Running a schedule

Issue the following command to run the schedule:

```
tpctool> runoptschedule -schedule_id 811002
```

The following output is returned:

```
Job ID
=====
18006
```

setarray

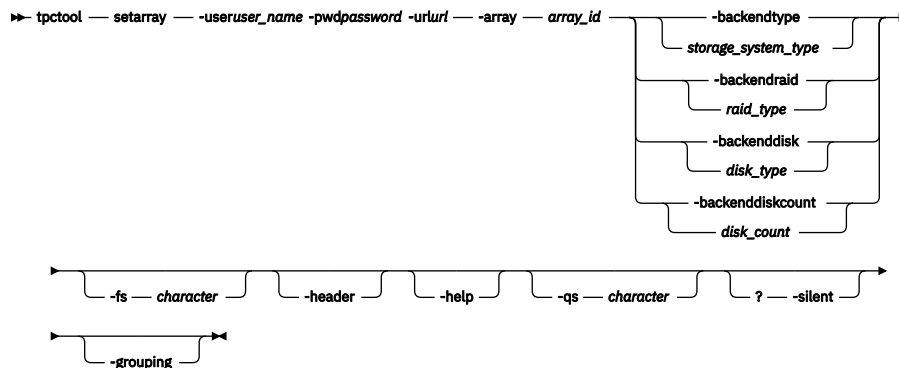
Use the **setarray** command to set the type of back-end storage system, type of Redundant Array of Independent Disks (RAID), type of disk, and number of disks for an array.

This command is available for the following storage systems:

- Storwize® V7000
- Storwize V7000 Unified
- SAN Volume Controller

You must have Administrator authority to use this command.

Syntax



Parameters and arguments

When you issue the **setarray** command, you must enter a value for the array parameter and at least one of the following parameters:

- -backendtype
- -backendraid
- -backenddisk
- -backenddiskcount

The default value is used for the parameters that you do not set.

Tip: You can set the values for back-end storage systems on the MDisk Group Details page in the IBM Spectrum® Control graphical user interface.

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- array array_id
Specifies the ID of the array.
- backendtype storage_system_type
Specifies the type of back-end storage system that manages most of the subsystem pool resources. You can enter a type that you added or use one of the following values.

Value	Name
C	EMC Clariion
D	DS8000®

Value	Name
S	EMC Symmetrix
X	XIV®
DEFAULT	Default value

Tip: Use the **lsbackendtypes** command to provide a list of existing back-end types of storage systems and use the **setbackendtype** command to add new back-end types of storage systems.

-backendraid raid_type

Specifies the type of RAID that is associated with the back-end storage system. You can enter a type that you added or use one of the following values.

Value	Description
1	RAID 1
5	RAID 5
6	RAID 6
X	RAID X
DEFAULT	Default value

Tip: Use the **lsbackendraidtypes** command to provide a list of existing back-end types of RAID and use the **setbackendraidtype** command to add new back-end types of RAID.

-backenddisk disk_type

Specifies the type of disk. You can enter a type that you added or use one of the following values.

Value	Description
A07	Sata - 7,500 rpm
F10	Fiber - 10,000 rpm
F15	Fiber - 15,000 rpm
DEFAULT	Default value

Tip: Use the **lsbackenddisktypes** command to provide a list of existing back-end types of disk and use the **setbackenddisktype** command to add new back-end types of disk.

-backenddiskcount disk_count

Specifies the number of disks.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

-grouping

Enables the grouping of numeric values. For example, in English the value 12000 would display as 12,000. The grouping character is determined by the system locale.

Example: Setting back-end storage subsystem values

The following command sets the back-end storage system, type of RAID, type of disk, and the number of disks for the specified array:

```
tpctool> setarray -array 0000020064405BA0:0+0000020064405BA0+0 -backendtype D
-backendraidtype 5 -backenddisktype F10 -backenddiskcount 160
```

The array ID and the status of the message is displayed.

```
Array List                               Status
=====
0000020064405BA0:0+0000020064405BA0+0 Succeeded.
```

setbackenddisktype

Use the **setbackenddisktype** command to set or update the type of back-end disk.


This command is available for the following storage systems:

- Storwize® V7000
- SAN Volume Controller


You must have Administrator authority to use this command.

Syntax for setting and updating back-end disk types

Use this syntax to set a new type of back-end disk:



Use this syntax to update the type of a back-end disk:



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- type disk_type
Specifies the type of the back-end disk that manages most of the back-end pool resources. The maximum number of characters that you can enter is 10. You can specify a back-end disk type or use one of the following values:

Value	Description
A07	Sata - 7,500 rpm
F10	Fiber - 10,000 rpm
F15	Fiber - 15,000 rpm
DEFAULT	Default type for unconfigured devices

- iops disk_iops
The average number of input/output operations per second for the disk. You can specify a value or use one of the following values.

IOPS Value	Back-end Disk Type
40	A07
120	F10
150	F15
0	Default value for unconfigured devices

- description description
Provides more information about the type of back-end disk. The maximum number of characters that you can enter is 256.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- grouping
Enables the grouping of numeric values. For example, in English the value 12000 would display as 12,000. The grouping character is determined by the system locale.

Example: Setting a back-end disk type

The following command sets the type of back-end disk for a back-end disk pool:

```
tpctool> setbackenddisktype -type testing -description testing -iops 99
```

If the command is successful, the following message is displayed:

```
Back-end Disk Type Status
=====
testing                Succeeded.
```

setbackendraidtype


Use the **setbackendraidtype** command to set the types of back-end RAID arrays for managed disk groups.

This command is available for the following storage systems:

- Storwize® V7000
- SAN Volume Controller

You must have Administrator authority to use this command.

Syntax



Parameters and arguments

-type *raid_type*

Specifies the type of the back-end RAID array. You can specify a back-end RAID type or use one of the following values:

Value	Description
DEFAULT	The default RAID type.
1	RAID 1
5	RAID 5
6	RAID 6
X	The RAID type that is used by IBM® XIV® Storage System.

Tip: Use the **lsbackendraidtypes** command to list the available types of back-end RAID types.

-weightedio *weighted_io*

Specifies the weighted input/output of the RAID type.

-user *user_name*

Specifies an IBM Spectrum Control user ID.

-pwd *password*

Specifies the password for the IBM Spectrum Control user ID.

-url *url*

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fs *character*

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

-help | -h | -?

Lists help information for the command.

-qs *character*

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

-grouping

Enables the grouping of numeric values. For example, in English the value 12000 would display as 12,000. The grouping character is determined by the system locale.

Example: Setting back-end RAID types

The following command sets a type of back-end RAID:

```
tptool> setbackendraidtype -type testing -weightedio 15
```

If the command is successful, the following message displays:

```
Back-end RAID Type Weighted IO
=====
testing              Succeeded.
```

setbackendtype

Use the **setbackendtype** command to set or update the type of back-end storage system.

This command is available for the following storage systems:

- Storwize® V7000
- Storwize V7000 Unified
- SAN Volume Controller

You must have Administrator authority to use this command.

Syntax for setting and updating back-end types of storage systems

Use this syntax to set a new type of back-end storage system:

```
tpctool --setbackendtype --user user_name --pwd password --url url --type storage_system_type --name storage_system_name --cachehit cache_hit_ratio --description description --fs character --header --help --qs character ? --silent --grouping
```

Use this syntax to update a type of back-end storage system:

```
tpctool --setbackendtype --user user_name --pwd password --url url --type storage_system_type --name storage_system_name --cachehit cache_hit_ratio --description description --fs character --header --help --qs character ? --silent --grouping
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- type storage_system_type
Specifies the type of back-end storage system that manages most of the subsystem pool resources. The maximum number of characters that you can enter is 10. You can enter a type or use one of the following values:

Value	Name
C	EMC Clariion
D	DS8000®
S	EMC Symmetrix
X	XIV®
DEFAULT	Default value

- name storage_system_name
Specifies the name of the back-end type of storage system. The maximum number of characters that you can enter is 256. The name of the type of storage system is displayed on the MDisk Group Details page in the IBM Spectrum® Control graphical user interface. You can enter a name or use one of the following values.

Description value	Associated type value
"EMC Clariion"	C
DS8000	D
"EMC Symmetrix"	S
XIV	X
DEFAULT	DEFAULT

Tip: If you enter a value for the -name or -description parameter that contains spaces, you must enclose the value in double quotation marks such as "My Description".

- cachehit cache_hit_ratio
Specifies the ratio of cache hits for read operations. The maximum number of characters that you can enter is 4. You can specify a value or use one of the following values:

Value	Back-end storage system
50	EMC Clariion
50	DS8000
70	EMC Symmetrix
50	XIV
0	Default value

- description description
Provides more information about the type of back-end storage system. The maximum number of characters that you can enter is 256. You can enter a description or use one of the following values:

Value	Description
"CX BE Cache Hit Ratio"	EMC Clariion
"DS8K Disk Controller"	DS8000
"DMX BE Cache Hit Ratio"	EMC Symmetrix
"IBM XIV"	XIV
"DEFAULT TYPE FOR UNCONFIGURED DEVICES"	0 (default value)

- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.

- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- grouping
Enables the grouping of numeric values. For example, in English the value 12000 would display as 12,000. The grouping character is determined by the system locale.

Example: Setting a new type of back-end storage system

The following command sets a new type of back-end storage system:

```
setbackendtype -type L -name test3 -cachehit 90 -description test_3
```

The following message is displayed:

```
Back-end Type Status
=====
L           Succeeded.
```

setdscfg

Use the **setdscfg** command to set the value of a property in the property file for the Device server. You must have Administrator authority to use this command.

Syntax

```
➤ tpctool — setdscfg — -user user_name -pwd password -url url — -property — property_key —>
                                     -help
? — ? — ? — ? — value — ? — ? — -silent —>
      -context context
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- help | -h | -?
Lists help information for the command.
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- property property_key value
Specifies the value to set for the indicated property key. The *property_key* variable is the property key and the *value* variable is the value.
- context context
Specifies a classification or category for a configuration property. The *context* variable is the context properties. For example, `-context DeviceServer` applies to the IBM Spectrum® Control Device server only. The parameter, `-context PerformanceManager`, applies to the IBM Spectrum Control performance manager only.

Example: Setting a property value

The following command sets the value of the **SnmpRetryCount** property to 3:


```
tpctool setdscfg -url localhost:9550 -user ***** -pwd ***** -property
SnmpRetryCount -context DeviceServer 3
```

setquota

Use the **setquota** command to set the amount of disk space and number of inodes that are assigned on a file system for a specified user name, group, or fileset.

A file system consists of attributes that include user names, groups, and filesets. You can set the amount of disk space and the number of inodes that are assigned for each attribute. This action limits the attributes from using more than the maximum amount of space or inodes that it was assigned. You must have Administrator authority to use this command.

Syntax



Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- usr | -group | -fileset
Specifies the type of quota.
- filesystem file_system_id
Specifies the ID for the file system that is associated with the quota. You can use the **lsfs** command to view information, including the IDs, for all file systems that are discovered. The ID is listed in the ID column of the **lsfs** command output.
- sspace s_space
Specifies the usage soft limit or level of disk space at a level less than that which the user, group, or fileset can safely operate. If you enter a size without a suffix, then the unit of measurement is byte. Otherwise, enter the suffix k (kilobyte), m (megabyte), g (gigabyte), t (terabyte), or p (petabyte).
- hspace h_space
Specifies the usage hard limit or maximum disk space which the user, group, or fileset can accumulate. If you enter a size without a suffix, then the unit of measurement is byte. Otherwise, enter the suffix k (kilobyte), m (megabyte), g (gigabyte), t (terabyte), or p (petabyte).
- sinode s_inode
Specifies the soft limit or the number of inodes space at a level less than that which the user, group, or fileset can safely operate. You can enter inode limits with only k (kilobyte), or m (megabyte) suffixes. The maximal value that you can specify is 2 GB.
- hinode h_inode
Specifies the hard limit or maximum number of inodes which the user, group, or fileset can accumulate. You can enter inode limits with only k (kilobyte), or m (megabyte) suffixes. The maximal value that you can specify is 2 GB.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.

Example: Set a quota for a user (-usr)

The following command sets the amount of disk space and number of inodes that are assigned on a file system for the user `tpcuser`:

```
tpctool> setquota -filesystem ee_fs03+kq98n5d.ibm+00000200A20045DC+0
-usr tpcuser -sspace 70M -hspace 100M -user admin -pwd password
-url localhost:9550
```

The following output is returned:

```
QuotaId          Status
=====
4971_U tpcuser+00000200A20045DC+0 SUCCESS
tpctool>
```

Example: Set a quota for a group

The following command sets the amount of disk space and number of inodes that are assigned on a file system for the group `users`:

```
tpctool> setquota -filesystem ee_fs03+kq98n5d.ibm+00000200A20045DC+0
-group users -sspace 50M -hspace 100M -user admin -pwd password
-url localhost:9550
```

The following output is returned:

```
QuotaId          Status
=====
4971_G_users+00000200A20045DC+0 SUCCESS
```

Example: Set a quota for a fileset

The following command sets the amount of disk space and number of inodes that are assigned on a file system for the fileset `testfset`:

```
tpctool> setquota -filesystem ee_fs03+kq98n5d.ibm+00000200A20045DC+0
-fileset testfs -sspace 50M -hspace 100M -user admin -pwd password
-url localhost:9550
```

The following output is returned:

QuotaId	Status
=====	=====
4971_F_testfs+00000200A20045DC+0	SUCCESS

Related reference

- [lsfs](#)

Related information

- http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp

showoptresults

Use the **showoptresults** command to show the recommendations for optimizing the placement of volumes.

In the web GUI, you use the Analyze Tiering wizard to specify criteria for analyzing the tiering of volumes. For example, you specify the following criteria:

- The storage pools that contain the volumes that you want to analyze
- The target storage pools for placing volumes that are underutilized and overutilized.
- The storage tiering policy to place volumes in the storage tier that best meets the performance requirements of the volumes

If the workload activity requirements of the analyzed volumes are not met, recommendations for optimizing the placement of the volumes are provided. In the web GUI, you can create schedules that are based on the criteria that you entered in the Analyze Tiering wizard. Each time the schedule is run, a job with a unique job ID is created.

Tip: To issue the **showoptresults** command, you must enter the job ID that is associated with the schedule. Issue the **lsoptschedules** command to show a list of the schedule IDs and their associated job IDs.

When you issue the **showoptresults** command, the following information is displayed for each volume.

Recommendation ID

Specifies the unique ID for the recommendations.

Volume

Specifies the name of the volume that is selected. For volumes in mirrored volume relationships, it specifies the name of the primary volume.

Source Pool

Specifies the current location of the volume.

Destination Pool

Specifies the storage pool that best meets the workload requirements of the volume.

Source Tier

Specifies the tier level of the source storage pool.

Destination Tier

Specifies the tier level of the destination storage pool.

Volume Copy Pool

Specifies the original location of the secondary volume of a volume in a mirrored volume relationship.

Server

Specifies the name of the server that the volume is assigned to.

Syntax

```
➔ showoptresults -job_id job_id -user user_name -pwd password -url url
                                     -fs character -header
                                     -help -qs character
```

Parameters and arguments

-job_id job_id

Specifies the ID of the job.

-user user_name

Specifies an IBM Spectrum Control user ID.

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-fs character

Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

-header

Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

- help | -h | -? Lists help information for the command.
- qs character Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").

Example: Show the analysis results for a job

Issue the following command to show the analysis results for a job:

```
tpctool> showoptresults -job_id 11002
```

The following output is returned:

Recommendation ID	Volume	Source Pool
68066	volume_1	pool_a
68067	volume_2	pool_c

Destination Pool	Source Tier	Destination Tier
pool_b	2	1
pool_f	1	3

Volume Copy Pool	Server
-	-
-	-

start

Use the **start** command to start a transaction. You must have Administrator authority to use this command.

Syntax

```
tpctool — start — -useruser_name -pwdpassword -urlurl — -fabric — WWN — -help
```

Parameters and arguments

- user user_name Specifies an IBM Spectrum Control user ID.
- pwd password Specifies the password for the IBM Spectrum Control user ID.
- url url Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- fabric WWN Specifies the fabric. The *WWN* variable is the worldwide name (WWN).
- help | -h | -? Lists help information for the command.

Example: Starting a transaction

The following command starts a transaction:

```
tpctool> start -user me -pwd mypass -url myhost:myport -fabric 100000051E34F6A8
```

unassignvol

Use the **unassignvol** command to remove the host ports from the assignment list for a volume. You must have Administrator authority to use this command.

Syntax

```
tpctool — unassignvol — -useruser_name -pwdpassword -urlurl — -hp — host_port — -f — -fs — character — -dev — GUID — -header — -help — -qs — character — ? — -silent — volume_id
```

Parameters and arguments

- user *user_name*
Specifies an IBM Spectrum Control user ID.
- pwd *password*
Specifies the password for the IBM Spectrum Control user ID.
- url *url*
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- hp *host_port*
Specifies the host ports. The *host_port* variable is a comma-separated list of worldwide port numbers (WWPNs).
- f
Turns off the confirmation message that is displayed before the ports are removed from the assignment list.
- fs *character*
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- dev *GUID*
Specifies the storage subsystem. The *GUID* variable is the globally unique identifier (GUID) of the storage subsystem as returned by the **lsdev -sys** command.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs *character*
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- volume_id | -
Specifies the volumes. The *volume_ID* variable is a comma-separated list of volume IDs, such as that obtained by running the **lsvol** command. If a single dash (-) is issued, the volume IDs are read from standard input.

Example: Removing host ports

The following command removes three host ports from the assignment list for the specified volume:

```
tpctool> unassignvol -user me -pwd mypass -url myhost:myport  
-hp 5005076300C79470,5005076300D09470,5005076300CB9470 2105.22232
```

The following output is returned:

Volume ID	PoolID	Status
2105.22232	P0	SUCCESS
2105.22232	P1	SUCCESS
2105.22232	P2	SUCCESS

unlinkfset

Use the **unlinkfset** command to unlink a fileset. You must have Administrator authority to use this command.

Unlinking a fileset makes all the files in the fileset inaccessible for as long as the fileset remains unlinked. The files are not deleted; they are only inaccessible.

Syntax

```
tpctool — unlinkfset — -user user_name -pwd password -url url — -fileset file_set_id — -f —  
-help —  
-header — ? — -silent — -qs — character — -fs — character
```

Parameters and arguments

- user *user_name*
Specifies an IBM Spectrum Control user ID.
- pwd *password*
Specifies the password for the IBM Spectrum Control user ID.
- url *url*
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- help | -h | -?
Lists help information for the command.
- fileset *file_set_id*
Specifies the IBM Spectrum® Control key of the fileset to be unlinked. The fileset key is listed in the **ID** column of the **lsfset** command output.
- f
Forces the files to be unlinked. If you use this option, the command forcibly closes any open files, causing the **ESTALE** error the next time the file is used.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.

- silent
Suppresses all output for the command. If you omit this parameter, output is included.
- qs character
Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks (").
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.

Example: Unlinking a fileset

The following command unlinks the fileset named `eefset01+eefs+kq458mv.ibm+00000200A2A0153C+0`.

```
tpctool> unlinkfset -fileset eefset01+eefs+kq458mv.ibm+00000200A2A0153C+0
-user admin -pwd password -url localhost:9550
```

The following output is returned:

FilesetId	Status
eefset01+eefs+kq458mv.ibm+00000200A2A0153C+0	SUCCESS

Related reference

- [rmfset](#)
- [lsfset](#)

Related information

- http://publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp

unmountfs

Use the **unmountfs** command to unmount a file system.

The **unmountfs** command unmounts a file system on all interface nodes, management nodes, or a specified subset. You must have Administrator authority to use this command.

Attention: Exports of the file system can block the unmount process. The active exports of the file system must be inactive before the unmount process is done.

Syntax

```
tpctool — unmountfs — -user user_name -pwd password -url url — -filesystem file_system_id — -nodes nodes —
- wait — -fs character — -header — -help — -qs character —
? — -silent —
```

Parameters and arguments

- user user_name
Specifies an IBM Spectrum Control user ID.
- pwd password
Specifies the password for the IBM Spectrum Control user ID.
- url url
Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.
- filesystem file_system_id
Specifies the ID of file system to be unmounted.
- nodes nodes
Lists the nodes to unmount the file system on, in a comma-separated list. Specify only the interface node, management node, or both. If you omit this parameter, the file system is unmounted on all nodes.
- wait
Indicates that the system waits until the file system is unmounted on all the nodes. An error occurs after the system waits 3 minutes.
- fs character
Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
Lists help information for the command.
- qs character

Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").

-silent

Suppresses all output for the command. If you omit this parameter, output is included.

Example: Unmount a file system

The following command unmounts a file system on a management node.

```
tpctool> unmountfs -filesystem eefs+kq458mv.ibm+00000200A2A0153C+0  
-user admin -pwd password -url localhost:9550 -nodes mgmt001st001
```

The following output is returned:

FilesystemId	Status
=====	
eefs+kq458mv.ibm+00000200A2A0153C+0	SUCCESS

Related information

- 🔗 <https://www.ibm.com/support/knowledgecenter/ST5Q4U/>

updatesrg

Use the **updatesrg** command to update the attributes of the specified Storage Resource Group with a new name, description, or user-defined properties.

Syntax

```
tpctool — updatesrg — -user user_name -pwd password -url url — -currentname currentSRGName — -newname →  
→ newSRGName — -newdescription newdescription — -newudp1 user-defined-property1 — -newudp2 →  
→ user-defined-property2 — -newudp3 user-defined-property3 — -fs character — -header →  
→ -help — -qs character — ? — -silent →
```

Parameters and arguments

-user user_name

Specifies an IBM Spectrum Control user ID.

-pwd password

Specifies the password for the IBM Spectrum Control user ID.

-url url

Specifies the Device server. The format of the URL is *system:port_number*, where *system* represents either the host name or IP address, and *port_number* represents the IBM Spectrum Control Device server port.

-currentname currentSRGName

Specifies a unique name already in use, fewer than 60 characters long. It cannot contain any of the following characters:

`\ / : * ? " < > | .`

The name of the user who creates the Storage Resource Group forms the beginning of the name. Use the **lssrg -l** command to display a list of Storage Resource Groups.

-newname newSRGName

Specifies a unique name not already in use, fewer than 60 characters long. It cannot contain any of the following characters:

`\ / : * ? " < > | .`

The name of the user that creates the Storage Resource Group is added to the beginning of the Storage Resource Group name, followed by the name that is specified in the **newSRGName** variable.

-newdescription newdescription

Specifies a description that can contain any alphanumeric characters, and must be fewer than 255 characters long.

-newudp1 user-defined-property1

(optional) Specifies a user-defined property that might contain anything of significance to the Storage Resource Group administrator. The UDP can contain any alphanumeric character, and must be fewer than 255 characters long.

-newudp2 user-defined-property2

(optional) Specifies a user-defined property that might contain anything of significance to the Storage Resource Group administrator. The UDP can contain any alphanumeric character, and must be fewer than 255 characters long.

-newudp3 user-defined-property3

(optional) Specifies a user-defined property that might contain anything of significance to the Storage Resource Group administrator. The UDP can contain any alphanumeric character, and must be fewer than 255 characters long.

-fs character

- Specifies the 7-bit character that separates the information in the output. If you omit this parameter, the information is separated by blank spaces.
- header
 - Suppresses the column headings in the output. If you omit this parameter, the column headings are included.
- help | -h | -?
 - Lists help information for the command.
- qs character
 - Specifies the 7-bit character that surrounds character strings and date stamps in the output. If you omit this parameter, character strings and date stamps are enclosed by double quotation marks ("").
- silent
 - Suppresses all output for the command. If you omit this parameter, output is included.

Example: Updating the attributes of a Storage Resource Group

The following command updates a Storage Resource Group with the specified attributes:

```
tpctool
tpctool> updatesrg -currentname currentSRGName -newname newSRGName
-newdescription newdescription -newudp1 user-defined-property1
-newudp2 user-defined-property2 -newudp3 user-defined-property3
```

The following output is returned:

Name	Status
Administrator.my-new-example-srg	SUCCESS

Command aliases

This topic discusses command aliasing.

With aliasing, you define a name for the alias followed by a value that is the name of a command and any options associated with command. The aliased command string is replaced by the defined value and the entire line is reparsed. Passwords used in aliased commands must first be encrypted using the **encrypt** command.

Aliased commands are saved in the command configuration file. The default configuration file is c:\program files\ibm\tpc\cli\libs\tpccli.conf.

For example, to shorten a frequently used command, you can define the following alias:

```
tpctool>lsperf = lsdev -user dsadmin -pwd 1ac75d82784ce0a327d45289604ae7b227
-url 9.44.33.126:8990 -fabric -perf
```

After the alias is defined, you can run the **lsperf** command to run the aliased **lsdev** command previously displayed.

You can provide a short form command targeting different device servers, as follows:

```
tpctool>lsperfd1 = lsdev -user dsadmin -pwd 1ac75d82784ce0a327d45289604ae7b227
-url hostOne:9161 -perf

tpctool>lsperfd2 = lsdev -user dsadmin -pwd 1ac75d82784ce0a327d45289604ae7b227
-url hostTwo:9161 -perf
```

You can specify additional options and arguments for an aliased command:

```
tpctool>lsperfd2 -fabric -ctype port
```

The command is expanded as follows:

```
tpctool>lsdev -user dsadmin -pwd 1ac75d82784ce0a327d45289604ae7b227
-url hostTwo:9161 -perf -fabric -ctype port
```

You can also nest aliases:

```
tpctool>lsperf = lsdev -user dsadmin -pwd 1ac75d82784ce0a327d45289604ae7b227
-url hostTwo:9161 -perf -fabric

tpctool>lsperfd1 = lsdev -user dsadmin -pwd 1ac75d82784ce0a327d45289604ae7b227
-url hostOne:9161

tpctool>lsperfd2 = lsdev -user dsadmin -pwd 1ac75d82784ce0a327d45289604ae7b227
-url hostTwo:9161
```

To unset an alias, type the name of the command alias followed by the '=' sign:

```
lsperf =
```

Parameter aliases

This topic lists common parameters and their aliases.

The following list shows parameters and their corresponding aliases.

- pwd
 - Password. The password is automatically encrypted using the same encryption algorithm as the **password** command before being stored in the config file. In conjunction with the interactive mode, this enables secure password encryption (plain text passwords will not appear in a command line).

-url
URL
-fs
Field-separator
-silent
Suppress-output
-header
Show-header
-l
Long
-dev
subsystem
-fabric
Fabric
-svr
Server

agent.sh command

The **agent.sh** command lets you start, stop, and restart the Storage Resource agent. You can also display the status and version of the Storage Resource agent.

Note:

- You must have root authority to run this command.
- For Windows, the status, stop, and start functions are handled through the Windows Services panel.

➔ **agent.sh** — status — start — stop — restart — version — versionall ➔

Parameters:

status
Displays the current status of the Storage Resource agent. The status is returned is running or not running.

start
Starts the Storage Resource agent.

stop
Stops the Storage Resource agent.

restart
Stop and then start the Storage Resource agent.

version
Displays the current version of the Storage Resource agent.

versionall
Displays the version of the Storage Resource agent and its related components.

dataCollector command

The dataCollector command is used to start and stop the data collector service.

About this task

The asset, capacity, and performance metadata for the storage systems in your data centers is collected by the data collector, analyzed, and then shown on the IBM Spectrum® Control GUI. The data collector service is part of the Device server. When you stop and start the Device server, the data collector stops and starts automatically.

Note:

Important: Do not use the dataCollector command unless you are directed by IBM® Support.

On Windows operating systems, you must have Administrator rights to stop or start the data collector service. On AIX® or Linux® operating systems, you must have root privileges to stop or start the data collector service.

Procedure

1. Log on to the server where the IBM Spectrum Control servers are installed.
2. Open a command window or shell script and go to the data collector directory.
The data collector service is installed in the *installation_dir/datacollector* directory.
3. Choose one of the following options to stop or start the data collector service:
 - a. On a Windows operating system, click the Start menu, enter `services.msc` and press Enter.
 - b. Select IBM Spectrum Control data collector and then stop or start it.

Or

 - a. Run the **dataCollector.bat** script with the stop or start parameter.

Or

 - a. On an AIX or Linux operating system, run the **dataCollector.sh** script with the stop or start parameter.

Configuration files

Use the parameters in IBM Spectrum® Control configuration files to help resolve problems.

The parameters in the configuration files are case-sensitive.

The default *installation_dir* for IBM Spectrum Control installations is as follows:

Windows operating systems:

c:\Program Files\IBM\TPC

Linux® or AIX® operating systems:

/opt/IBM/TPC

The default file locations for the configuration files for IBM Spectrum Control are as follows:

IBM Spectrum Control:

Windows operating systems:

installation_dir\config

Linux or AIX operating systems:

installation_dir/config

Data server:

Windows operating systems:

installation_dir\data\config

Linux or AIX operating systems:

installation_dir/data/config

Device server:

Windows operating systems:

installation_dir\device\conf

Linux or AIX operating systems:

installation_dir/device/conf

Storage Resource agent:

Windows operating systems:

installation_dir\agent\config\agent.config

Linux or AIX operating systems: operating systems:

installation_dir/agent/config/agent.config

Restriction: On Windows installations, if you installed IBM Spectrum Control by using a domain user account, you must disable User Account Control to edit the configuration files.

- [server.config file](#)

The following information lists the parameters that are set in the server.config file. These parameters include controller, logging, repository, and service.

- [scheduler.config file](#)

The following information lists the parameters that are set in the scheduler.config file. These parameters include concurrency parameters and jobs parameters.

- [TPCD.config file](#)

The list of parameters that are set in the TPCD.config file include server parameters and GUI parameters.

- [Specifying the tablespace size for IBM Spectrum Control](#)

This section provides information on the size of the tablespace to specify when you install IBM Spectrum Control

- [agent.config file](#)

The agent.config file contains configuration parameters for the Storage Resource agent. These parameters are set when the Storage Resource agent is installed. The parameters can also be changed manually by editing the file.

server.config file

The following information lists the parameters that are set in the server.config file. These parameters include controller, logging, repository, and service.

Controller parameters

name

The Data Manager server name is the name of the host computer.

port

The port on which the server listens for requests. The default is 9549.

maxConnections

The maximum number of concurrent sockets that the server opens. The default is 500.

routerThreads

The number of threads that redirect incoming requests to the appropriate service provider. The default is 1.

serviceThreads

The number of threads to allocate for the internal service provider of the server. The default is 2.

agentErrorLimit

The number of consecutive attempts to reach an agent before the agent is displayed as DOWN. The default is 3. When an agent is in this state, no attempts to connect are made until either the agent contacts the server or the agent status is manually changed to UP.

adminGroup

The name of the group a user must be a member of to perform administrative functions from the GUI, the default is **adm**.

commEncrypted

The switch that secures communication between the Server or Agent and the Server/GUI by encrypting the DataStream.

- **0** = Off. Do not encrypt the DataStream.
- **1** = On. Encrypt the DataStream.

FileSystemScan NFSTimeout

Determines the numbers of seconds that a Storage Resource agent waits for a status system call on a Network File System (NFS) before it times out.

hostAlias

This parameter is displayed if the HOST_ALIAS is not specific and represents the name of the server. The value for this parameter is used when multiple computers have the same name or the name cannot be determined.

Logging parameters

logsKept

The number of server logs to keep. The default is 5.

messagesPerLog

The maximum number of messages in a log. When this number is reached the log is closed and a new log is created. The default is 100,000.

Repository parameters

driver

The name of the JDBC driver to use, normally:

- Db2®: **COM.ibm.db2.jdbc.app.DB2Driver**

url

The URL used to connect to the database, normally:

- Db2: **jdbc:db2:database_name**

user

The user name that IBM Spectrum® Control uses to connect to the repository.

connectionPool

The number of database connections in a pool of reusable open connections. The default is 10.

Service parameters

name

Repeating section that indicates the service providers that are required to start.
The REQUIRED parameters are as follows:

- **TStorm.server.svp.GuiSvp**
- **TStorm.server.svp.AgentSvp**
- **scheduler.Scheduler**

scheduler.config file

The following information lists the parameters that are set in the scheduler.config file. These parameters include concurrency parameters and jobs parameters.

Concurrency parameters

maxSubmitThreads

The maximum number of threads to create that handle the submission of jobs. The default is 3.

maxCompleteThreads

The maximum number of threads to create to handle job completions. Initially creates a pool of half the number of threads specified that can grow to the maximum.
The default is 3.

Jobs parameters

minutesAdvanced

The number of minutes in advance of scheduled time to begin the scheduling process. The default is 1. Use this option to allow for the processor time that is involved in scheduling a job so that the job starts close to the scheduled time.

delayLimitMinutes

Number of minutes after scheduled start time that the Scheduler continues to attempt to start a job for a selected resource, so that resource state is not in a down state or, connection status is not in a failed state. The default is 120.

Location of the scheduler.config file

The scheduler.config file is in the following directories:

Windows operating systems:
installation_dir\data\config

Linux® or AIX® operating systems:
installation_dir\data/config

TPCD.config file

The list of parameters that are set in the TPCD.config file include server parameters and GUI parameters.

The following list describes the server parameters:

threadPoolSize

Number of initial threads to create for handling requests. The default is 3.

abbreviatedProbe

Only SCSI commands are sent to disk drives for inquiry and disk capacity information. The default is 1.

maxThreads

Set the maximum number of threads for handling requests. The default is 8.

pingReceiveTimeout

Number of seconds to wait before it indicates that a ping failed. The default is 10.

skipAutoFS

Set to 1 if you want to skip the **automount** process during discovery on the Oracle Solaris Storage Resource agent. By default, discovery always processes **automount** on all Oracle Solaris Storage Resource agents managed by the Data server.

saveNonRoot

Set to 1 if you want to monitor non root exports. The default is 0.

If you do not set this parameter, the export paths that are not at the root of the file system are discarded. If the NAS server has only non root exports accessible to the agent, it will not be added. Restart the Data server for this setting to take effect.

batchPartitionWaitRetryCount

Specify the number of times that the Storage Resource agent tries to get a report partition before the Storage Resource agent fails with an error.

Large batch reports are generated in partitions. The partitions are placed on the IBM Spectrum® Control server, and the Storage Resource agent gets them from the server when the batch report is created.

The following list describes the GUI parameters:

threadPoolSize

Number of initial threads to create for handling user interface requests. The default is 3.

maxThreads

Set the maximum number of threads for handling user interface requests. The default is 10.

reportRowLimit

Maximum number of rows that are sent at a time to the user interface. If this number is exceeded, a More button is displayed over the table, along with a warning message. The default is 5000.

keepCachedReport

Number of minutes to retain incomplete reports in the tmp directory for the server. The default is 120.

Specifying the tablespace size for IBM Spectrum Control

This section provides information on the size of the tablespace to specify when you install IBM Spectrum® Control

About this task

When you install IBM Spectrum Control, you can specify the tablespace size of the repository database or accept the default values. The space needed for the IBM Spectrum Control database schema varies significantly with storage network configuration, data collection, data retention period, and other factors.

The following table provides space estimates for a storage configuration containing 5000 volumes with some general assumptions.

Table 1. Tablespace allocation for the IBM Spectrum Control database schema

Tablespace	Description of tablespace usage	Recommended size for a 5000 volume configuration	Assumptions
KEY	This tablespace is used for configuration data which is constantly used. For example, the key entity and relationships data (T_RES_STORAGE_SUBSYSTEM, T_RES_STORAGE_VOLUME, and the normalization tables, and so forth)	500 MB	

Tablespace	Description of tablespace usage	Recommended size for a 5000 volume configuration	Assumptions
NORMAL	This tablespace is used for snapshots and miscellaneous data	500 MB	A table that uses significant space is T_RES_STORAGE_VOLUME_SNAPSHOT. This table uses about 2500 bytes for each record. The number of snapshots depends on the data collection activities.
BIG	This tablespace is used for performance statistics	2 to 3 GB or 400* MB per day of performance data	The data collected for performance data for storage volumes can use a significant amount of space (about 200 bytes for each record). For 5000 volumes, if performance data is collected every 5 minutes, the data for one day would be 300 MB. If the data is kept for 7 days, the data collected would take about 2 to 3 GBs. If the data is kept longer, the storage must be scaled up accordingly.
TEMP	This tablespace is used for temporary data for query processing and other temporary tables	1GB	

agent.config file

The agent.config file contains configuration parameters for the Storage Resource agent. These parameters are set when the Storage Resource agent is installed. The parameters can also be changed manually by editing the file.

The following list contains the parameters for the agent.config file.

Servename

Fully qualified host name of the system on which the Data server is installed.

Portnumber

Port on which the Data server listens for communications from the Storage Resource agent. By default, the port is set to 9549.

IPAddress

IP address of the server on which the Data server is installed.

Log files

When you have a problem, you can check several product log files.

- [Default locations of log files](#)

Check the log files to view detailed information about IBM Spectrum Control processing and to troubleshoot problems.

Default locations of log files

Check the log files to view detailed information about IBM Spectrum® Control processing and to troubleshoot problems.

The following list shows the default log file locations for IBM Spectrum Control and other components.

Device server

The IBM® WebSphere® Liberty Profile log files for the Device server are in the following directories:

Windows operating systems.

`installation_dir\wlp\usr\servers\deviceServer\logs`

Linux® or AIX® operating systems.

`installation_dir/wlp/usr/servers/deviceServer/logs`

The operational log files for the Device server are in the following directories:

Windows operating systems.

`installation_dir\device\log`

Linux or AIX operating systems.

`installation_dir/device/log`

The log files for the data collector are in the following directories:

Windows operating systems.

`installation_dir\datacollector\log`

Linux or AIX operating systems.

`installation_dir/datacollector/log`

Alert server:

The IBM WebSphere Liberty Profile log files for the Alert server are in the following directories:

Windows operating systems.

`installation_dir\wlp\usr\servers>alertServer\logs`

For example, C:\Program Files\IBM\TPC\wlp\usr\servers>alertServer\logs

Linux or AIX operating systems.

`installation_dir/wlp/usr/servers/alertServer/logs`

The operational log files for the Alert server are in the following directories:

Windows operating systems.

installation_dir\alert\log

For example, C:\Program Files\IBM\TPC>alert\log

Linux or AIX operating systems.

installation_dir/alert/log

Data server

Windows operating systems.

installation_dir\data\log

Linux or AIX operating systems.

installation_dir/data/log

Export server

Windows operating systems.

installation_dir\export\logs

Linux or AIX operating systems.

installation_dir/export/logs

Web server log files

Windows operating systems.

installation_dir\wlp\usr\servers\webServer\logs

Linux or AIX operating systems.

installation_dir/wlp/usr/servers/webServer/logs

IBM Spectrum Control GUI

Windows operating systems.

installation_dir\web\log

Linux or AIX operating systems.

installation_dir/web/log

Storage Resource agents

installation_dir/agent/log/name_of_server_SRA_communicates_with

Tips:

- For Windows operating systems, the default *installation_dir* is C:\Program Files\IBM\TPC.
- For Linux or AIX operating systems, the default *installation_dir* is /opt/IBM/TPC.

Script parameters

Script parameters provide specific information on the alert that triggered the script to be run.

The parameters that are passed to a script depend on the type of alert that was triggered. The following table describes all the script parameters:

Script Parameter	Description
amount	Threshold exceeded amount.
archive-file-count	The number of log files in the archived log directory.
archive-log-directory	Name of the archive log directory that triggered the archive log directory Instance alert.
available-extents	The number of extents still available to the segment for growth. This value equals the maximum extents available to the object minus the extents that are currently allocated to the segment.
available-space	Available pool space after a change
blade	Name of a blade.
chained-row-count	The number of chained rows in a table that triggered the Chained Row table alert.
computer	Computer name where the triggering condition occurred.
consecutive-failures	Number of consecutive failed attempts to ping the computer.
controller	Name of a back-end controller.
cluster-name	The name of an HACMP or MSCS cluster.
creator.name	Creator of the ping, probe, or scan schedule. Name of the schedule.
current-grown-defects	Current® number of grown defects on the disk.
current-node-name	When an HACMP or MSCS cluster resource group moves, this parameter identifies the cluster node that now hosts the cluster resource group.
current-RAM MB	Current value of the RAM in megabytes.
current-VM MB	Current value of the sum of the RAM and the swap space in megabytes.
database	The name of the database where the triggering condition occurred.
database-tablespace	The name of the database or table space where the triggering condition occurred.
device-name	Name of a device.
disk-array	Name/alias of a disk array.
dump-date	The date when the last memory dump was performed.
endpoint	Name of an endpoint device.
extent-count	The number of extents that are allocated to a segment, or the number of free extents in the table space (depends on Alert type).
failed-jobs	Number of failed jobs in the run. (Each job runs on a different computer).
file-of-violating-files	Temporary file that contains a list of files that violate the constraint. The files are listed as one file per line.

Script Parameter	Description
file-of-violating-owners	Temporary file that contains a list of owners who owns the violating files.
free-inodes	Maximum number of files available to be created on this file system.
free space size-designator	Total amount of free space, in KB, MB, or GB.
from-entity-type	Type of new fabric connection from an entity.
HBA-driver	HBA driver
HBA-firmware	HBA firmware
io-group	Name of the I/O group.
largest-extent-size size-designator	Total amount of the largest free extent in the table space, which is measured in KB, MB, or GB.
manufacturer/serial-number	Manufacturer of the disk. Serial number of the disk.
mdisk	Name of an MDisk.
mdisk-group	Name of an MDisk group.
mount-point	Path to the file system.
new-capacity	New capacity of a storage subsystem, volume, or pool.
new-version	New version of the HBA driver, firmware, or a subsystem.
node	Name of a node.
old-capacity	Previous capacity of a storage subsystem, volume, or pool.
old-grown-defects	Previous number of grown defects on the disk.
old-node-name	When an HACMP or MSCS cluster resource group moves, this parameter identifies the cluster node that previously hosted the cluster resource group.
old-RAM MB	Previous value of the RAM in megabytes.
old-version	Previous version of the HBA driver, firmware, or a subsystem.
old VM MB	Previous value of the sum of the RAM and the swap space in megabytes.
path	Path to the directory.
percent-of-capacity %	Percentage of capacity of the file system, database, or table space.
percent-of-table-size	The percentage of space that is allocated to a segment that is empty and unused (the percentage of space over the "high-water mark"). Available on the Empty Used Segment Space table alert.
percent-of-total-rows %	The percentage of table rows that are chained.
pool	Name of a storage pool.
port	Name of a port.
rdbms-instance-name	Oracle SID, SQL Server name, Sybase Server name, UDB Instance name
rdbms-type	Oracle, SQL Server, or Sybase
run-number	Number of the run.
segment	The name of the table segment that triggers the alert.
segment-type	<p>The type of segment that triggers the alert. The following list includes the possible types of segments.</p> <ul style="list-style-type: none"> • TABLE • TABLE PARTITION • TABLE SUBPARTITION • NESTED TABLE • CLUSTER • INDEX • INDEX PARTITION • INDEX SUBPARTITION • LOBINDEX • LOBSEGMENT • LOB PARTITION • LOB SUBPARTITION
storage-volume	Name of a storage volume
subsystem	Name of a storage subsystem
switch	Name of a switch
table	The name of the table that triggered the alert condition.
table space	The name of the table space that triggered the alert condition.
threshold	Value that you set for the triggering condition. If the value unit was specified as a %, then a % follows this value.
threshold thr-designator	Value of the triggering condition, in KB, MB, or GB, or % (value units).
to-entity-type	Type of new fabric connection to an entity.
total-jobs	Total number of jobs in a run.
total-file-size size designator	Total amount of storage that is consumed by the archive log directory, which is measured in KB, MB, or GB.
usage size-designator	Value of used disk space, in KB, MB, or GB.
violating-file-count	Number of files that met the conditions that are defined in the constraint.
virtual-server-name	The name of an HACMP or MSCS cluster resource group.
zone	Name of a zone.
zoneset	Name of a zone set.
zone-alias	Name of a zone alias
zone-member	Name of a zone member

Related reference

- [How scripts are run](#)

Opening IBM Spectrum Control on Windows operating systems

You can open IBM Spectrum Control CLIs and GUIs and administer IBM Spectrum Control on Windows operating systems.

- [Opening IBM Spectrum Control GUIs and CLIs](#)
To manage and monitor storage resources, open IBM Spectrum Control GUIs and CLIs.
- [Accessing administration tools](#)
To manage and maintain IBM Spectrum Control, access the Windows system administration tools.

Opening IBM Spectrum® Control GUIs and CLIs

To manage and monitor storage resources, open IBM Spectrum Control GUIs and CLIs.

About this task

You can open the following GUIs and CLIs:

- [Opening IBM Spectrum Control GUI](#)
- [Opening Db2 Command Window](#)
- [Opening IBM Data Studio Administration client on Windows operating systems](#)
- [Opening IBM Tivoli Monitoring Services](#)

Opening IBM Spectrum Control GUI

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	a. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start. b. On the Start page, right-click, and then click the All apps taskbar button.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start, > All Programs.

2. Click IBM Spectrum Control > IBM Spectrum Control.

Opening Db2® Command Window

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	a. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start. b. On the Start page, right-click, and then click the All apps taskbar button.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start, > All Programs.

2. Click IBM DB2 > Command Line Tools > Command Window.

Opening IBM® Data Studio Administration client on Windows operating systems

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	a. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start. b. On the Start page, right-click, and then click the All apps taskbar button.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start, > All Programs.

2. Click IBM Data Studio > Data Studio Administration Client.

Opening IBM Tivoli® Monitoring Services

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	a. On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start. b. On the Start page, right-click, and then click the All apps taskbar button.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start > All Programs.

2. Click IBM Tivoli Monitoring > IBM Tivoli Monitoring Services.

Accessing administration tools

To manage and maintain IBM Spectrum Control, access the Windows system administration tools.

About this task

To complete tasks in IBM Spectrum Control, you must open the following administration and maintenance facilities:

- [Accessing the Control Panel](#)
- [Accessing Administrative Tools](#)
- [Accessing Windows Services](#)
- [Accessing Computer Management](#)
- [Accessing Programs and Program Features](#)

Accessing the Control Panel

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start.

2. Click Control Panel

Accessing Administrative Tools

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start.

2. Click Administrative Tools

Accessing Windows Services

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start.

2. Click Administrative Tools > Services

Accessing Computer Management

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start.

2. Click Administrative Tools > Computer Management

Accessing Programs and Program Features

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start.

2. Click Control Panel > Programs and Program Features.

Accessing Window Run

Procedure

1. Choose one of these options:

Option	Description
Windows Server 2012	On the Dashboard page, hover the mouse over the lower left corner of the page next to the Server Manager taskbar button, and then click Start.
Windows Server 2012 R2, Windows Server 2016, Windows Server 2019	Click Start.

2. Click Run.

Windows services used by IBM Spectrum Control

To start, stop, or restart a component or related program in IBM Spectrum® Control, use the Windows Services panel.

The following table provides a list of Windows services.

Table 1. List of Windows services used by IBM Spectrum Control

Program	Service name	Comment
IBM® DB2®	DB2 - DB2COPY1 - DB2 - 0 DB2 Governer (DB2COPY1) DB2 License Server (DB2COPY1) DB2 Management Service (DB2COPY1) DB2 Remote Command Server (DB2COPY1) DB2DAS - DB2DAS00 DB2TS - DB2COPY1 - DB2-0	The service account owner is db2admin . The account needs to be part of Administrators and DB2ADMNS.
IBM Spectrum Control Data server	IBM Spectrum Control - Data Server	
IBM Spectrum Control Device server	IBM Spectrum Control - Device Server	
IBM Spectrum Control Alert server	IBM Spectrum Control - Alert Server	
IBM Spectrum Control Storage Resource agent	IBM Spectrum Control Storage Resource Agent	
IBM Spectrum Control Web server	IBM Spectrum Control - Web Server	
IBM Spectrum Control Export server	IBM Spectrum Control - Export Server	
IBM Spectrum Control data collector	IBM Spectrum Control data collector	

Frequently Asked Questions

View answers to common questions about IBM Spectrum® Control.

How do you know if your storage system is supported by IBM Spectrum Control and which SMI-S agents are supported?

To confirm whether your storage system is supported and which SMI-S agent is supported for that system, review the supported products list for the current release of IBM Spectrum Control at [IBM Spectrum Control interoperability matrix for storage systems](#).

You encounter errors while collecting performance data on SAN Volume Controller. You fail to associate SAN Volume Controller performance data from non-configuration node with SAN Volume Controller performance data from configuration node. You encounter incomplete SAN Volume Controller performance data sample.

This issue is caused by a configuration issue with SAN Volume Controller (time zone). Reset the time zone on SAN Volume Controller by logging into the SAN Volume Controller through putty. Run this command first:

```
svctask settimezone -timezone 509
```

This forces the cluster into the Universal time zone. To get the time zone you want the cluster to be in, run this command:

```
svctask settimezone -timezone
```

Protocols and standards

This section provides an overview of the protocols and standards that are used within IBM Spectrum® Control.

- **[Web Based Enterprise Management](#)**
Web Based Enterprise Management (WBEM) is an initiative of the Distributed Management Task Force (DMTF) with the objective to enable the management of complex IT environments. It defines a set of management and internet standard technologies in order to unify the management of complex IT environments.
- **[Storage Management Initiative Specification](#)**
The Storage Networking Industry Association (SNIA) defines a standard that is used within IBM Spectrum Control to create and develop a universal open interface for managing storage devices including storage networks.
- **[Service Location Protocol](#)**
The Service Location Protocol (SLP) is an Internet Engineering Task Force (IETF) standard. SLP provides a scalable framework for the discovery and selection of network services.
- **[Simple Network Management Protocol](#)**
The Simple Network Management Protocol (SNMP) is an Internet Engineering Task Force (IETF) protocol for monitoring and managing systems and devices in a network. Functions supported by the SNMP protocol are the request and retrieval of data, the setting or writing of data, and traps that signal the occurrence of events.
- **[Fibre Channel Methodologies of Interconnects](#)**
IBM Spectrum Control supports the ANSI T11 Fibre Channel FC-MI (Fibre Channel Methodologies of Interconnects) for the automated discovery of FC SAN assets and topology.

Web Based Enterprise Management

Web Based Enterprise Management (WBEM) is an initiative of the Distributed Management Task Force (DMTF) with the objective to enable the management of complex IT environments. It defines a set of management and internet standard technologies in order to unify the management of complex IT environments.

The WBEM initiative is composed of three main conceptual elements:

Common Interface Model (CIM)

CIM is a formal object-oriented modeling language that is used to describe the management aspects of systems.

xmlCIM

This is the grammar to describe CIM declarations and messages used by the CIM protocol.

Hypertext Transfer Protocol (HTTP)

HTTP is used as a way to enable communication between a management application and a device that both use CIM.

The WBEM architecture defines the following elements:

CIM Client

The CIM Client is a management application like IBM Spectrum® Control that uses CIM to manage devices. A CIM Client can reside anywhere in the network, because it uses HTTP to talk to CIM Object Managers and Agents.

CIM Managed Object

A Managed Object is a hardware or software component that can be managed by a management application by using CIM.

CIM Agent

A CIM Object Manager that includes the provider service for a limited set of resources. An agent may be embedded or hosted and can be an aggregator for multiple devices.

CIM Provider

A CIM Provider is the element that translates CIM calls to the device-specific commands. A provider is always closely linked to a CIM.

CIM Object Manager (CIMOM)

The central component of the CIM Server responsible for the communication between the CIM server components.

CIM Server

A server that receives and processes CIM Operation Message Requests and issues CIM Operation Message Responses.

Storage Management Initiative Specification

The Storage Networking Industry Association (SNIA) defines a standard that is used within IBM Spectrum® Control to create and develop a universal open interface for managing storage devices including storage networks.

For information about SMI-S, see <http://www.snia.org>.

SNIA has fully adopted and enhanced the Common Information Model (CIM) standard for storage management in its Storage Management Initiative - Specification (SMI-S). SMI-S was launched to create and develop a universal open interface for managing storage devices including storage networks. SMI-S provides:

- A comprehensive specification for the management of heterogeneous storage and storage area networks (SANs).
- The information available to a WBEM client from an SMI-S compliant CIM server (provider).
- Profiles organized by:
 - Storage
 - Fabric
 - Host
 - Common profiles and subprofiles
- An object-oriented CIM and XML-based interface for managing SAN devices, services, and fabrics.
- An initial discovery, which is SLP based.

The idea behind SMI-S is to standardize the management interfaces so that management applications can utilize these and provide cross-device management. This means that a newly introduced device can be immediately managed as it will conform to the standards.

The models and protocols in the SMI-S implementation are platform-independent, enabling application development for any platform, and enabling them to run on different platforms. The SNIA will also provide interoperability tests which will help vendors test their applications and devices if they conform to the standard.

Service Location Protocol

The Service Location Protocol (SLP) is an Internet Engineering Task Force (IETF) standard. SLP provides a scalable framework for the discovery and selection of network services.

The Internet Engineering Task Force (IETF) is a large open international community of network designers, operators, vendors, and researchers that are concerned with the evolution of the Internet architecture and the smooth operation of the Internet. The IETF includes formal standards for SNMP and MIBs. For more information about IETF, see <http://www.ietf.org>.

SLP enables the discovery and selection of generic services, which can range in function from hardware services such as those for printers or fax machines, to software services such as those for file servers, email servers, web servers, databases, or any other possible services that are accessible through an IP network.

Traditionally, to use a particular service, a user, or client application provided the host name or network IP address for the service. With SLP, however, it is not necessary for the user or client application to know individual host names or IP addresses. Instead, the user or client application can search the network for the required service type and an optional set of qualifying attributes.

For example, a user can search for all available printers that support Adobe PostScript. Based on the service type such as printers and the attributes such as PostScript, SLP searches the user's network for matching services, and returns the discovered list to the user.

Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) is an Internet Engineering Task Force (IETF) protocol for monitoring and managing systems and devices in a network. Functions supported by the SNMP protocol are the request and retrieval of data, the setting or writing of data, and traps that signal the occurrence of events.

SNMP is a protocol that enables a management application to query information from a managed device. The managed device has software running that sends and receives the SNMP information. This software module is usually called the SNMP agent.

An SNMP management application can read information from an SNMP agent in order to monitor the device that the SNMP agent is running on. Therefore, the device needs to be polled on an interval bases. The SNMP manager can also change the configuration of a device, by setting certain values to corresponding variables. A device can also be set up to send a notification to the SNMP manager (this is called a trap) to asynchronously inform this SNMP manager of a status change.

Depending on the existing environment and organization it is very likely that your environment already has an SNMP management application in place. IBM Spectrum® Control can be set up to send traps.

SNMP uses a hierarchical structured Management Information Base (MIB) to define the meaning and the type of a particular value. A MIB defines managed objects that describe the behavior of the SNMP entity, which can be anything from an IP router to a storage subsystem. The information is organized in a tree structure.

For users planning to make use of the IBM Spectrum Control SNMP trap alert notification capabilities, an SNMP MIB is included in the server installation.

The MIB is provided for use by your SNMP management console software. Most SNMP management products provide a program called a MIB compiler that can be used to import MIBs. This will allow you to view IBM Spectrum Control generated SNMP traps from within your management console software. Refer to your management console software documentation for instructions on how to compile or import a third party MIB.

For a Cisco switch to successfully receive and respond to queries from IBM Spectrum Control, the following basic requirements must be met:

- IBM Spectrum Control can use SNMPv3 (preferred) or SNMPv1 to probe switches and fabrics. The SNMPv3 protocol is preferred because it provides better security, but switches that use the SNMPv1 protocol are also supported. Some switches are configured to use SNMPv3 by default.
- The Fibre Alliance FC Management MIB (FA MIB) and Fibre Channel Fabric Element MIB (FE MIB) must be enabled on the switch.
- When you use the SNMPv1 protocol, the community string that is configured in IBM Spectrum Control must match one of the community strings that are configured on the switch with read access. Additionally, Cisco switches must have a community string match for write access. The default community strings in IBM Spectrum Control are "public" for read access and "private" for write access. Other community strings can be defined on the switches, but are not used. Community strings are not relevant when you use the SNMPv3 protocol.
- SNMP access control lists must include the IBM Spectrum Control system. These access control lists are defined and configured on the switches. Some lists automatically include all hosts, while others exclude all by default.
- The Fibre Channel (FC) or Fibre Channel over Ethernet (FCoE) protocols must be enabled on the switch. Some switches, such as the Cisco Nexus 5000 series, require you to enable these protocols. Otherwise, IBM Spectrum Control does not recognize the switch when you try to add it using the Add Switches and Fabrics for Monitoring dialog. For instructions on how to configure Cisco switches for FCoE enablement, go to the Cisco product website at <http://www.cisco.com> and click Support.

IBM Spectrum Control uses port 162 to listen for SNMP traps. This is the default port. For switches, you must configure the switch to send SNMP traps to the Device server IP address. If you need to change the default port number, use the **setdscfg** command. The attribute to set is **SNMPTrapPort**.

System administrators must set up their SNMP trap ringer with the provided MIB files in order to receive SNMP traps from IBM Spectrum Control. These files are located in the following directories on the product installation DVD:

For the Data server:

`data\snmp\tivoliSRM.mib`

For the Device server :

`device\snmp\fabric.mib`

Fibre Channel Methodologies of Interconnects

IBM Spectrum® Control supports the ANSI T11 Fibre Channel FC-MI (Fibre Channel Methodologies of Interconnects) for the automated discovery of FC SAN assets and topology.

ANSI T11 Fibre Channel FC-MI includes the following for the automated discovery of FC SAN assets and topology:

- Hosts (HBAs)
- FC interconnects
- FC storage devices

The T11 FC-MI also includes the following:

- FC-GS-3/4 (discovery, zoning, and so forth)
- RNID (advanced device recognition)
- Platform registration (device recognition and launch)
- Common HBA API (fabric and storage views)
- Name server (connectivity)
- Management server (SAN connectivity and topology)
- RSCN (advanced event detection)
- SCSI queries (storage views, volume information, and so forth)
- SNMP Fabric Element (FE) MIB
- SNMP FC Management MIB (discovery, performance statistics, and so forth)
- SNMP alerts

IBM Spectrum Control technical community

Connect, learn, and share with storage professionals: product support technical experts who provide their perspectives and expertise.

Access the IBM Spectrum® Control technical community at <https://www.ibm.com/developerworks/servicemanagement/>.

Use IBM Spectrum Control technical community in the following ways:

- Become involved with transparent development, an ongoing, open engagement between other users and developers of IBM® products.
- Connect one-on-one with the experts to collaborate and network about IBM and the Storage Management community.
- Read blogs to benefit from the expertise and experience of others.
- Use forums to collaborate with the broader user community.

Accessibility features for IBM Spectrum Control

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in IBM Spectrum® Control:

- Keyboard-only operation in the GUI.
- The IBM® Documentation that includes the following accessibility features:
 - The IBM Documentation is provided in XHTML 1.0 format, which is viewable in most web browsers. With XHTML, you can view documentation according to the display preferences that are set in your browser. XHTML supports screen readers and other assistive technologies.
 - All documentation for IBM Spectrum Control is available in Adobe Portable Document Format (PDF) by using the Adobe Acrobat Reader. To access PDFs, go to [Printable documentation](#).
 - All images in the IBM Documentation are provided with alternative text, so that visually impaired users can understand the contents of the images.
- Interfaces that are commonly used by screen readers.

The setting for the automatic-refresh button in the screen reader is toggled to the ON position by default. If you want the screen reader to read the previous text, complete the following steps:

1. Navigate to the Accessibility Settings Navigation region by using the arrow keys. The region is located after the IBM Spectrum Control application title.
2. Click Enter to toggle the automatic-refresh button to the OFF position. (An alert sounds to make you aware that the turn-off automatic-refresh toggle button was pressed.)
3. To move backward to the previously read text so that the screen reader can read it again, use the arrow keys. You can move backwards and forwards through the page.
4. When you are ready to move on, click Enter to toggle the automatic-refresh button to the ON position and to refresh the page. (An alert sounds to make you aware that the turn-on, automatic-refresh toggle button was pressed.)

Tip: Alternatively, let the toggle setting persist, and refresh as needed by pressing the F5 key.

Keyboard navigation

Most of the features of the IBM Spectrum Control GUI are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI), except as noted in the product release notes.

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. The following sections describe the keys or key combinations for different parts of the GUI:

For navigating in the GUI and the context-sensitive help system:

- To navigate to the next link, button, or topic within a panel, press Tab.
- To move to the previous link, button, or topic within a panel, press Shift+Tab.
- To select an object, when the object is in focus, press Enter.

For actions menus:

- To navigate to the grid header, press Tab.
- To reach the drop-down field, press the Left Arrow or Right Arrow key.
- To open the drop-down menu, press Enter.
- To select the menu items, press the Up Arrow or Down Arrow key.
- To start the action, press Enter.

For filters:

To specify a filter option and text:

1. Press Tab to navigate to the magnifying glass icon.
2. Press the Up Arrow or Down Arrow key to navigate the filtering list.
3. Press Enter to select a filtering option.
4. When a filtering option is selected, the cursor moves to the filter text box. Type the filter text and press Enter. To reset a filter, press Enter.

For text fields:

- To navigate to text fields, press Tab.
- To navigate to the fields that are available for editing, press Tab.
- To navigate to the next field or to the Submit button, press Tab.

For tables or lists:

- To navigate between column headers, focus on a column header and use the Left Arrow and Right Arrow keys to move to other column headers.
- To navigate between data cells, focus on a data cell and use the Left, Right, Up, Down, Pageup, and Pagedown Arrow keys.
- To sort a column, focus on a column header and press Enter. The focus remains on the column header after the sort occurs.
- To change the size of a column, focus on the column header, hold Shift+Control, and press the Left or Right Arrow keys.
- To follow a link in a data cell, focus on a data cell and press Shift+F9.
- To open a menu for a table row, focus on the row and press Shift+F10.
- To select consecutive rows, select the first row and hold Shift, press the Up or Down Arrow keys to go to the last row in the range, and press the Space bar to add the new rows to the selection.
- To select non-consecutive rows, select a row and hold Control, press the Up or Down Arrow keys, and press the Space bar to add the new row to the selection.

Restriction: For Chinese languages, the keyboard combination Control+Space bar is not enabled for selecting multiple rows at the same time.

Keyboard navigation with Firefox for Mac users: If you're using Firefox on a Mac with IBM Spectrum Control and want to use keyboard navigation, complete the following steps:

1. In Firefox, go to Preferences, > Advanced, > General and clear the check mark for Always use the cursor keys to navigate within pages. This step enables the use of Tab key to navigate between GUI elements.
2. In the URL address bar of Firefox, type about:config and press Enter.
Tip: If a warning prompt is displayed, click the button to accept the risk of changing browser settings. Existing settings won't be changed; instead, you'll be adding a preference setting for accessibility.
3. To add an accessibility preference for tab focus, right-click on the configuration page and select New, > Integer.
4. In the New integer value window, type accessibility.tabfocus and click OK.
5. Type 7 to set the integer value and click OK.
6. Open your Mac's System Preferences app, go to Keyboard, > Shortcuts, and select All Controls.

IBM and accessibility

For more information about IBM's commitment to accessibility, see the IBM Human Ability and Accessibility Center website at <http://www.ibm.com/able>.

Troubleshooting and problem determination

When a problem occurs with IBM Spectrum® Control, you want to find a solution quickly. You can resolve a problem by using two approaches: troubleshooting and problem determination.

Troubleshooting

In many cases, you can resolve problems without contacting the IBM® Support Center. When you encounter a problem with IBM Spectrum Control, you can search the troubleshooting section of the Knowledge documentation for a solution. Troubleshooting information explains the cause of a problem at a high level and describes the actions that you can take to resolve the problem.

Problem determination

If you cannot resolve a problem by using the troubleshooting information that is provided in the IBM Spectrum Control Knowledge documentation, contact the IBM Support Center so that the cause of your problem can be determined. To learn how to gather information for the IBM Support Center, use the problem determination section of the Knowledge documentation.

- [Troubleshooting in IBM Spectrum Control](#)
You can use this information to resolve issues or problems with IBM Spectrum Control.
- [Getting support](#)
For help with resolving issues with IBM Spectrum Control, you can contact IBM Support or use IBM self-help resources.

Troubleshooting in IBM Spectrum Control

You can use this information to resolve issues or problems with IBM Spectrum® Control.

- [Introduction to troubleshooting](#)
The troubleshooting information explains the cause of a problem at a high level and describe actions that you can take to resolve the problem.
- [Resolving problems](#)
If you experience an error with IBM Spectrum Control, you can use problem determination to determine why an error occurred, and explain how to resolve the problem.

Introduction to troubleshooting

The troubleshooting information explains the cause of a problem at a high level and describe actions that you can take to resolve the problem.

- [GUI troubleshooting](#)
Troubleshoot and resolve problems in the IBM Spectrum Control GUI.
- [General troubleshooting](#)
Use this section to troubleshoot and resolve problems with the IBM Spectrum Control GUI.
- [Db2 and database troubleshooting](#)
Use this section to troubleshoot and resolve problems with IBM® Db2® and the database.
- [Installation, uninstallation, and upgrading](#)
Use this section to troubleshoot and resolve problems with IBM Spectrum Control installation, uninstallation, and upgrading.
- [SAN Volume Controller troubleshooting](#)
Use this section to troubleshoot and resolve SAN Volume Controller problems.
- [DS8000 troubleshooting](#)
Use this section to troubleshoot and resolve DS8000® problems.
- [Independent software vendors troubleshooting](#)
Use this section to troubleshoot and resolve independent software vendor problems.
- [Language troubleshooting](#)
Learn how to troubleshoot and resolve problems with non-English versions of the product.

GUI troubleshooting

Troubleshoot and resolve problems in the IBM Spectrum® Control GUI.

- [Renewing security certificates](#)
You can renew the IBM Spectrum Control server certificates that are expiring or have expired.
- [Displaying performance metadata](#)
Occasionally, you might find that performance metadata cannot be displayed in IBM Spectrum Control. This issue might occur when metadata is not available for the selected time range or device, or when the cache, cookies, and history in your web browser need to be cleared to refresh the GUI with the latest metadata.
- [Troubleshooting FTP transfers](#)
When you attempt to send files to IBM® by using FTP, such as a log package from the System Management page or a compressed file from the service tool, the upload cannot be completed. This problem might occur because of network or firewall restrictions.
- [The status of jobs is not updated after the Device server is stopped](#)
The statuses of data collection jobs in the IBM Spectrum Control GUI are not updated after the Device server is stopped.
- [Changing the session timeout value for the web-based GUI](#)
A session of the web-based GUI times out and returns to the login page.
- [Error message: This web browser is not supported](#)
This message is displayed on the login page in Internet Explorer 8.
- [The GUI installation program displays the incorrect title on AIX 7.x with the UTF-8 locale](#)
This problem is cosmetic only and does not impact the installation of IBM Spectrum Control on the AIX® 7.x operating system when you use UTF-8 locales.
- [Error message: Out of memory at line: 14](#)
If you use Internet Explorer, the details page of the GUI might become unresponsive after the page is open for an hour or more.
- [Interface icons are not displayed in the correct positions](#)
The icons on the dashboard, resource lists, and details pages in the GUI are not displayed in the correct positions and do not respond to mouse or keyboard actions.
- [The status of a virtual machine is not consistent](#)
The status of a monitored virtual machine might be different depending on the page in the GUI where it is being viewed.
- [Information exported to a CSV file is incorrectly formatted by spreadsheet program](#)
Information that is displayed in the web-based GUI can be exported to a comma-separated values (CSV) file. When you are viewing a CSV file in a spreadsheet program, some information might not match what is displayed in the web-based GUI.

Renewing security certificates

You can renew the IBM Spectrum® Control server certificates that are expiring or have expired.

Before you begin

Note: This procedure assumes that you install IBM Spectrum Control on a Windows operating system by using the default installation paths and ports.
Make a backup copy of the following files:

- `installation_dir\wlp\usr\servers\deviceServer\resources\security\key.p12`
- `installation_dir\wlp\usr\servers>alertServer\resources\security\key.p12`
- `installation_dir\wlp\usr\servers\deviceServer\server.xml`
- `installation_dir\wlp\usr\servers>alertServer\server.xml`
- `installation_dir\wlp\usr\servers\webServer\resources\security\key.p12`
- `installation_dir\wlp\usr\servers\webServer\server.xml`

About this task

When secure socket layer (SSL) certificates expire or are close to expiring, you need to renew them. If you continue to receive error messages about the web server certificate expiration, use this more advanced method.

Procedure

1. Stop the IBM Spectrum Control servers and services.
2. Delete the following files: `installation_dir\wlp\usr\servers\deviceServer\resources\security\key.p12`; `installation_dir\wlp\usr\servers>alertServer\resources\security\key.p12`; `installation_dir\wlp\usr\servers\webServer\resources\security\key.p12`.
3. Open a command window and run the following commands:

- a. `cd installation_dir\wlp\bin`
- b. `securityUtility.bat createSSLCertificate --server=deviceServer --password=default --validity=825`
- c. `securityUtility.bat createSSLCertificate --server=alertServer --password=default --validity=825`
- d. `securityUtility.bat createSSLCertificate --server=webServer --password=default --validity=825`

For each of the preceding commands, you see output similar to the following example:

```
Creating keystore
C:\Program Files\IBM\TPC\wlp\usr\servers\deviceServer\
resources\security\key.p12

Created SSL certificate for server deviceServer

Add the following lines to the server.xml to enable SSL:
<featureManager> <feature>ssl-1.0</feature>
</featureManager>
<keyStore
id="_home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_troubleshooting_fqz0_t_tbs_renewing_security_cert
ificates_defaultKeyStore" password="{xor}Ozo5PiozKw=" />
```

Note: Do not add the example lines to the server.xml files. Those lines are already included in the server.xml files.

- e. `cd installation_dir\jre\bin`
 - f. `ikeycmd -cert -delete -db installation_dir\wlp\usr\servers\deviceServer\resources\security\key.p12 -pw default -type pkcs12 -label default`
 - g. `ikeycmd -cert -create -db installation_dir\wlp\usr\servers\deviceServer\resources\security\key.p12 -pw default -label default -size 2048 -sig_alg SHA256_WITH_RSA -expire 825 -dn "CN=<machine_FQDN>, OU=deviceServer, O=ibm, C=us" -san_dnsname <machine_FQDN> -eku serverAuth`
 - h. `ikeycmd -cert -delete -db installation_dir\wlp\usr\servers>alertServer\resources\security\key.p12 -pw default -type pkcs12 -label default`
 - i. `ikeycmd -cert -create -db installation_dir\wlp\usr\servers>alertServer\resources\security\key.p12 -pw default -label default -size 2048 -sig_alg SHA256_WITH_RSA -expire 825 -dn "CN=<machine_FQDN>, OU=alertServer, O=ibm, C=us" -san_dnsname <machine_FQDN> -eku serverAuth`
 - j. `ikeycmd -cert -delete -db installation_dir\wlp\usr\servers\webServer\resources\security\key.p12 -pw default -type pkcs12 -label default`
 - k. `ikeycmd -cert -create -db installation_dir\wlp\usr\servers\webServer\resources\security\key.p12 -pw default -label default -size 2048 -sig_alg SHA256_WITH_RSA -expire 825 -dn "CN=<machine_FQDN>, OU=webServer, O=ibm, C=us" -san_dnsname <machine_FQDN> -eku serverAuth`
- Where `<machine_FQDN>` is the fully qualified domain name of the machine where you installed the IBM Spectrum Control servers. For example, `myserver.mycompany.com`
4. If you renewed the Device server certificate, complete steps 4, 5, and 6 in *Updating IBM Spectrum Control data collector trusted certificates after replacing default SSL certificate for the Device server*.
 5. Start the IBM Spectrum Control servers and services.
 6. Start the IBM Spectrum Control GUI and click Home, > System Management.
 7. Verify that the servers and services are running properly.

Results

The certificate is renewed.

What to do next

Verify the IBM Spectrum Control web server certificate renewal. To verify that the certificate is renewed, open the IBM Spectrum Control GUI. If the new certificate was successfully applied, no web browser certificate or security warnings are displayed. To view the web server certificate, complete the following steps. These steps might be different depending on your browser version.

Mozilla Firefox:

1. Right-click a GUI page and select View Page Info.
2. On the Page information page, click Security, > View Certificate.

Internet Explorer:

1. Right-click a GUI page and select Properties.
2. On the Properties page, click Certificates.

Related tasks

- [Stopping the IBM Spectrum Control servers by using the GUI](#)
- [Updating IBM Spectrum Control data collector trusted certificates after replacing default SSL certificate for the Device server with a self-signed certificate OR an external certificate](#)

Displaying performance metadata

Occasionally, you might find that performance metadata cannot be displayed in IBM Spectrum® Control. This issue might occur when metadata is not available for the selected time range or device, or when the cache, cookies, and history in your web browser need to be cleared to refresh the GUI with the latest metadata.

About this task

If performance metadata cannot be displayed in IBM Spectrum Control, you can try clearing the cache, cookies, and history in your web browser to refresh the GUI with the latest metadata.

Procedure

1. Clear your web browser's cache, cookies, and history using the following steps:

In Mozilla Firefox:	In Google Chrome:
a. Click Tools > Options and select Privacy & Security. b. Under Cookies and Site Data, click Clear Data. c. Under History, click Clear History.	a. Click Settings and select Privacy and security. b. Click Clear browsing data, accept the default choices, and click Clear data.

2. After the cache, cookies, and history is cleared, click the browser refresh button to update the product GUI.

For more detailed and up-to-date instructions about how to clear the cache, cookies, and history in different types and versions of web browsers, refer to the documentation for those browsers.

Troubleshooting FTP transfers

When you attempt to send files to IBM® by using FTP, such as a log package from the System Management page or a compressed file from the service tool, the upload cannot be completed. This problem might occur because of network or firewall restrictions.

About this task

When IBM Spectrum® Control uploads a log package to IBM Software Support, it is sent to the Enhanced Customer Data Repository (ECuRep) by using standard FTP protocol. The upload is secure and the data is encrypted with AES-128 key for each file part and RSA-2048 key for the key exchange. The name of the FTP server that receives the upload is ftp.ecurep.ibm.com.

If IBM Spectrum Control cannot connect to the FTP server because of firewall restrictions in your organization, you must open the necessary ports for FTP transfers from the IBM Spectrum Control server. Because standard, unencrypted FTP in passive mode is used for uploads (encryption is done at the data level, not the connection level), the firewall must be configured to allow passive FTP connections to ftp.ecurep.ibm.com. Contact your network administrator for assistance.

For older firewalls, open port 21 (the default FTP) and ports 65024 - 65535 (for passive FTP). For more information, see http://www.ibm.com/de/support/ecurep/send_ftp.html.

If your organization uses a proxy server for FTP, complete the following steps to troubleshoot uploads from IBM Spectrum Control:

Procedure

1. Use a text editor to create a file that is named ibmsdduu.config.
2. In ibmsdduu.config, specify the proxy configuration to use.
 - a. To determine the specific information that must be included in ibmsdduu.config for your FTP proxy, review the following instructions:

SOCKS Server and HTTP Proxy Support

```
-socks4 enable SOCKS 4 support
-socks5 enable SOCKS 5 support
-http_proxy enable support for http proxy tunneling

-socks4 -sock5 and -http_proxy are exclusive, the last switch in
```

command line is used

-proxyhost=<host> define the address of the proxy or SOCKS server

-proxyport=<port> define the port of the proxy.
default: 1080 for SOCKS4 and SOCKS5
8080 for http_proxy

-proxyuser=<userid> define the userid for the proxy or SOCKS server

-proxypw=<password> define the password for the proxy or SOCKS server

proxy userid and proxy password are optional

FTP Proxies

ftp proxy support is enabled with a separate command line option.
This enables the combination of SOCKS server and a ftp proxy if necessary.

-ftp_proxytype=<type> enables ftp proxy support and define the type of
the proxy. <type> is an integer, see
ftp proxy types below
Default: 0 - no proxy

-ftp_proxyuser userid for the ftp proxy

-ftp_proxypw password for the ftp proxy

-ftp_proxyhost hostname or ip address of the ftp proxy

-ftp_proxyport port number of the ftp proxy
Default: 21

FTP Proxy Types

There are several ftp proxy types known. Each type is using a different
login procedure and needs different commands to connect to the target
ftp server.

In the list of supported ftp proxy types <server>,<user>,<password>
are referring to the target ftp server, <p_user> and <p_passwd> to the
ftp proxy account.

-ftp_proxytype=0
no ftp proxy is used (this is the default)

-ftp_proxytype=1
connect; USER <p_user>; PASS <p_passwd>;
SITE <server>; USER <user>; PASS <passwd>

-ftp_proxytype=2
connect; SITE <server>; USER <user>; PASS <passwd>

-ftp_proxytype=3
connect; USER <p_user>; PASS <p_passwd>;
OPEN <server>; USER <user>; PASS <passwd>

-ftp_proxytype=4
connect; OPEN <server>; USER <user>; PASS <passwd>

-ftp_proxytype=5
connect; USER <user>@<server>; PASS <passwd>

-ftp_proxytype=6
connect; USER <p_user>@<server>; PASS <p_passwd>;
USER <user>; PASS <passwd>

-ftp_proxytype=7
connect; USER <user>@<server> <p_user>;PASS <p_passwd>;
ACCT <passwd>

-ftp_proxytype=8
connect; USER <user>@<server> <p_user>;PASS <passwd>;
ACCT <p_passwd>

-ftp_proxytype=9
connect; USER <user>@<p_user>@<server>;PASS <passwd>;
ACCT <p_passwd>

-ftp_proxytype=10
connect; USER <user>@<p_user>@<server>;PASS <p_passwd>;
ACCT <passwd>

-ftp_proxytype=11
connect; USER <p_user>; PASS <p_passwd>;
USER <user>@<server>; PASS <passwd>

-ftp_proxytype=12 (CheckPoint Firewall 1)
connect; USER <user>@<p_user>@<server>
PASS <passwd>@<p_passwd>

- b. Based on the type of FTP proxy server that you use, copy and paste the appropriate section from step 2a into ibmsdduu.config.
- c. For the section that you pasted into ibmsdduu.config, provide information such as host name, port, user ID, and password, as needed.

For example:

```
-ftp_proxytype=1
-ftp_proxyuser=my_user_id
-ftp_proxypw=my_password
-ftp_proxyhost=ftp_proxy.my_company.com
-ftp_proxyport=21
```

Where:

- **my_user_id** and **my_password** represent the authentication credentials for connecting to the FTP proxy server.
- **ftp_proxy.my_company.com** represents the DNS name of the FTP proxy server.

3. Save the file to *installation_dir/services*, where *installation_dir* represents the directory where IBM Spectrum Control is installed.

The default installation directory is as follows:

- Windows: C:\Program Files\IBM\TPC
- Linux® / AIX®: /opt/IBM/TPC

4. Try to upload the log package or compressed file again.

What to do next

If you still cannot automatically upload a log package in the IBM Spectrum Control GUI, you can try uploading it manually. For information about how upload files to ECuRep manually, see the following topics:

- [FTP Transfer](#)
- [Upload through browser](#)

Log packages are stored on the IBM Spectrum Control server at the following location: *installation_dir/wlp/usr/servers/webServer/apps/WebServer.ear/TPC-GUI.war/serviceLog*

The status of jobs is not updated after the Device server is stopped

The statuses of data collection jobs in the IBM Spectrum® Control GUI are not updated after the Device server is stopped.

Problem

If the Device server is stopped while data collection jobs are running, the jobs are also stopped. However, even though the jobs are no longer collecting data, their statuses are still displayed as "running" in the GUI.

Action

To display the correct status of the jobs in the GUI, restart the Device server. For information about how to start the Device server, see [Starting the IBM Spectrum Control servers by using scripts](#).

Changing the session timeout value for the web-based GUI

A session of the web-based GUI times out and returns to the logon page.

Problem

The Lightweight Third Party Authentication (LTPA) token timeout value that is set in the IBM® WebSphere® Application Server Liberty ltpa.xml file determines when you are automatically logged out of the GUI. The default timeout value is 120 minutes.

Action

To change the timeout value for the LTPA token, complete the following steps:

1. Log on to the IBM Spectrum® Control server with administrative privileges.
2. Open a command prompt and go to *installation_dir/wlp/usr/servers/webServer* directory.
3. Edit the ltpa.xml file and change the value of the *expiration* attribute from 120 to the desired value.
4. Save and exit the ltpa.xml file.

Error message: This web browser is not supported

This message is displayed on the logon page in Internet Explorer 8.

Problem

You cannot log on because Internet Explorer 8 is not supported with this version of the product.

Action

To ensure that you can log on to IBM Spectrum® Control, upgrade to Internet Explorer 9 or later. For a list of web browsers that you can use with IBM Spectrum Control, see [IBM Spectrum Control interoperability matrix](#). In the Agents, Servers and Browser column, click the version of IBM Spectrum Control that is installed on your system. On the next page, click Web Browsers to find the web browsers that you can use.

The GUI installation program displays the incorrect title on AIX 7.x with the UTF-8 locale

This problem is cosmetic only and does not impact the installation of IBM Spectrum® Control on the AIX® 7.x operating system when you use UTF-8 locales.

Problem

The IBM Spectrum Control installation program displays the incorrect title on the AIX 7.x operating system when you use a UTF-8 locale identifier.

Action

For UTF-8 locales, if you installed IBM Spectrum Control on an AIX 7.x operating system and set your locale identifier to a UTF-8 locale (for example, fr_FR.UTF-8), the window title in the installation program is incorrect. The incorrect title is displayed as "sun-awt-X11-XFramePeer".

For non-UTF-8 locales, if you installed IBM Spectrum Control on an AIX 7.x operating system and set your locale identifier to a non-UTF-8 locale (for example, FR_FR), the window title in the installation program is correct. The correct title is displayed as "IBM Spectrum Control".

Error message: Out of memory at line: 14

If you use Internet Explorer, the details page of the GUI might become unresponsive after the page is open for an hour or more.

Problem

When you use Internet Explorer, if the details page of the GUI is open for long periods, the following message is displayed:

Out of memory at line: 14

Environment

This problem occurs on Windows operating systems when you use Internet Explorer.

Resolving the problem

Use Mozilla Firefox 3.6 or later to view the GUI. Alternatively, increase the memory of the system that is running Internet Explorer, or restart Internet Explorer and log on to the GUI again.

Related reference

- [Web browser support](#)

Interface icons are not displayed in the correct positions

The icons on the dashboard, resource lists, and details pages in the GUI are not displayed in the correct positions and do not respond to mouse or keyboard actions.

Problem

The GUI does not display correctly in Internet Explorer 9 when Compatibility View is turned on. When Compatibility View is on, the pages in the GUI are displayed as if you were using an earlier version of Internet Explorer.

Action

To ensure that the GUI is displayed correctly in Internet Explorer 9, complete the following steps:

1. In Internet Explorer 9, go to Tools > Compatibility View Settings.
2. In the section Websites you've added to Compatibility View, ensure that the website address for the GUI is not included.
3. Clear the following check boxes:
 - Include updated website lists from Microsoft
 - Display intranet sites in Compatibility View
 - Display all websites in Compatibility View
4. Click Close to apply the changes to the web browser.
5. Refresh the current page to apply the changes to the GUI.

The status of a virtual machine is not consistent

The status of a monitored virtual machine might be different depending on the page in the GUI where it is being viewed.

Problem

For example, if a virtual machine that is managed by a hypervisor is down or unreachable, the following statuses might be displayed:

- On the details page for a hypervisor, the status of the virtual machine might be Normal.
- On the details page for a server, the status of the virtual machine might be Unreachable.

This situation occurs because different probes are collecting data about the status of a virtual machine:

- Probes of the managing hypervisor collect information about the status of the virtual machine configuration.
- Probes of the virtual machine collect information about the virtual computer that is hosted by the virtual machine.

Action

Based on the status that you want to view, go to the corresponding location in the GUI:

- For the status of the virtual machine configuration, go to the details page for a hypervisor.
- For the status of the virtual computer that is hosted by the virtual machine, go to the details page for a server.

Information exported to a CSV file is incorrectly formatted by spreadsheet program

Information that is displayed in the web-based GUI can be exported to a comma-separated values (CSV) file. When you are viewing a CSV file in a spreadsheet program, some information might not match what is displayed in the web-based GUI.

Problem

A CSV file is a plain-text file. No type information is provided for the values in the CSV file, but a spreadsheet program, such as Lotus® Symphony® or Microsoft Excel, can determine a type based on the value itself. If a value contains only numeric characters, a spreadsheet program might determine that it is a numeric value. It might also format the value in a particular style, such as in decimal format or scientific notation. Sometimes, this automatic reformatting of values results in values that are incorrectly displayed. For example, a worldwide port name (WWPN) should be treated as a text even if it contains only numeric characters. If a spreadsheet program converts the WWPN to a new format, it no longer reflects that actual name.

Action

If values in the CSV file do not match the values that are displayed when the file is opened in a spreadsheet program, consult the documentation for the spreadsheet program you are using. When you open a CSV file in spreadsheet program, you can typically override the automatic formatting behavior and specify how values are imported. To verify that the values in the CSV file match the values that are shown in the spreadsheet program, you can open the CSV file in a text editor.

General troubleshooting

Use this section to troubleshoot and resolve problems with the IBM Spectrum® Control GUI.

- [Alerts aren't being generated](#)
Alerts aren't being generated because the Alert server stopped responding.
- [Performance information is not displayed for a resource](#)
Performance information is not displayed for a resource in performance views and charts.
- [The Data server is shut down automatically](#)
If IBM Spectrum Control cannot reach the database repository after 30 minutes, the Data server is shut down automatically.
- [Fabric probe of DCFM CIM/OMs returns Java "Out of Memory" errors](#)
During a fabric probe of a DCFM CIM/OM, if Java™ "Out of Memory" errors are returned, use the following steps to increase the server memory allocation for the DCFM CIM/OM.
- [Cluster resource group alerts are not triggered](#)
This problem can occur on systems running PowerHA® SystemMirror® for AIX® or Microsoft Cluster Server (MSCS).
- [Slow performance when moving data from the database repository](#)
Repocopy and other DB2® movement utilities might perform slowly or be in a wait state when moving data from the IBM Spectrum Control database repository.
- [Error message when running repocopy on server that uses remote database](#)
This error occurs when you run **repocopy** on a server machine that uses a remote database.
- [The Data Server service fails with a logon failure when restarted](#)
This error occurs when using Window's Active Directory domain policy.
- [Specifying a LUN ID for the assignvol command](#)
This error occurs when specifying a LUN ID for the **assignvol** command.
- [Cannot monitor an EMC CIMOM on a Solaris server](#)
This condition occurs when you try to monitor a Pegasus-based CIMOM like the EMC CIMOM.
- [SQLCODE-440 displayed if install IBM Spectrum Control on system with bad system clock](#)
This error occurs when you install IBM Spectrum Control on a system with a clock that is set for a future time and not the current time.
- [CIM agent runs slowly](#)
This error occurs for the CIM agent.

- [Assigned LUN is not recognized by the host](#)
A LUN assigned to a host using IBM Spectrum Control is not recognized by the host, and IBM Spectrum Control does not report the error.
- [On AIX systems, numbers in the Storage Resource agent registry are printed as formatted numbers](#)
On AIX systems, numbers in the Storage Resource agent or configuration file are printed as formatted numbers.
- [Fabric probe job failed](#)
The fabric probe job failed.
- [Error Messages: Get SQLCODE: -964, SQLSTATE: 57011, Unexpected error occurred while performing \[a function\]](#)
You might receive these error messages when performing tasks within IBM Spectrum Control.
- [Data server crashes on AIX](#)
This error occurs when you are running the Data server on AIX.
- [No performance data is retrieved](#)
This error occurs when you try to get performance data.
- [Performance monitoring job fails immediately after it starts](#)
This error occurs when you run a performance monitoring job and it fails immediately after it starts.
- [Error message "performance data files could not be correlated"](#)
This error occurs when you uninstall the agent using the IBM Spectrum Control uninstallation program.
- [Performance correlation step takes a long time](#)
This error occurs when a performance correlation step is done.
- [Cannot define a performance monitor for a switch](#)
This error occurs when trying to define a performance monitor for a switch.
- [File system storing database tables runs out of space](#)
This problem occurs on systems running AIX 5.x.
- [Error message: There is a problem with this website's security certificate](#)
You receive this error message on Internet Explorer when you launch the DS8000® Element Manager.
- [Cannot discover fabric on Red Hat Linux using Emulex HBA](#)
A fabric probe operation performed on Red Hat® Linux® using a Storage Resource agent cannot find the Emulex HBA.
- [Cannot get version of Storage Resource agent on SUSE Linux](#)
You cannot get the version of the Storage Resource agent on SUSE Linux.
- [Cannot log on to two instances of IBM Spectrum Control by using the same browser](#)
You cannot log on to two instances of IBM Spectrum Control by using the same browser.
- [Warning in Probe Log for Storage Resource agents on AIX \(STA0115W\)](#)
If certain services, such as `mountd` daemon, are not running on the AIX Storage Resource agents, warning entries are displayed in the Probe log for Storage Resource agents deployed on AIX.
- [A generic Storage Resource agent probe error might be a registration problem / GUID mismatch](#)
A generic Storage Resource agent (SRA) probe error might occur when sending data to the server.
- [Following IBM Spectrum Control server upgrade, the Storage Resource agents \(SRA\) might lack the Upgrade action](#)
Immediately following an upgrade of the IBM Spectrum Control server, the SRAs will have the Agent State of *Upgrade needed*, but might not have the option to start the SRA upgrade.
- [How to fix IBM Storwize V7000 Unified authentication errors](#)
If the IP address for IBM® Storwize® V7000 Unified is not configured correctly on SAN Volume Controller, you get an authentication error when you add the storage system in IBM Spectrum Control.
- [Slow performance on VMware virtual machines due to shared VMDK](#)
The VMware virtual machine where you run IBM Spectrum Control exhibits slow performance.
- [How to disable cipher block chaining \(CBC\) ciphers in IBM Spectrum Control in response to the Lucky 13 vulnerability](#)
If needed, you can disable cipher block chaining (CBC) ciphers in IBM Spectrum Control in response to the Lucky 13 vulnerability.

Alerts aren't being generated

Alerts aren't being generated because the Alert server stopped responding.

Problem

In some cases, the Alert server might be locked into a processing loop and stop responding when it processes performance alerts that include suppression settings. During this time, no alerts are generated or notifications sent for any alert definitions.

Action

To avoid this issue, complete the following steps:

1. In a text editor, open `installation_dir/bin/etc/Datastore.properties`.
2. Change `san.logger.trace.level=DEBUG_MAX` to `san.logger.trace.level=DEBUG_MIN`.
3. To apply the change, complete the following steps to restart the Alert server:
 - a. In the GUI, go to Home > System Management.
 - b. Click Component Servers in the Components section.
 - c. Click Stop Server next to the Alert server.
 - d. Click Start Server next to the Alert server.

Performance information is not displayed for a resource

Performance information is not displayed for a resource in performance views and charts.

Problem

For some resources and metrics, performance information is not displayed during a specific time range and resolution (granularity).

Action

If performance data is not being displayed for a resource and metric, one of the following situations might have occurred:

- A performance monitor was never defined for the resource. You must run a performance monitor to collect and view performance data. To schedule a performance monitor, go to the resource list page, right-click a resource, and select **Data Collection** > **Schedule**.
- A performance monitor was defined for the resource, but is not running. Accordingly, depending on when the most recent performance monitor completed, try adjusting the time range and data resolution on the performance view. You might need to change the data resolution to hourly or daily. Alternatively, start the performance monitor so that it can collect new data. To start a performance monitor, go to the resource list page, right-click a resource, and select **Data Collection** > **Start Performance Monitor**.
- A performance monitor was defined and is running, but cannot collect any data. Check the values of **Status** and **Latest Error** (message) for the performance monitor on the **Home** > **Performance Monitors** page. To resolve some common problems, try the following actions:
 - Verify that the local area network is available and a firewall is not preventing network access to the resource and product services.
 - Verify that the resource you want to monitor is up and available.
 - If an SMI-S provider (CIM agent) is being used to communicate with the resource, ensure that the provider can be accessed, and is up and available.
- None of the selected metrics is available. The metrics that are available for a resource are determined by a number of factors, such as the type of resource, release version of a resource, and level of SMI-S provider (if used). Select different metrics to view performance data.
- An internal component of Spectrum Control is not available or is down. To investigate the issue, try the following actions:
 - Check the status of the product servers and database repository on the **Home** > **System Management** page. In **Component Servers**, ensure that the servers are up and running. In **Database**, ensure that the database repository is up and running.
 - Check the log files that were generated for the internal components. [Location of log files](#).
- For FlashSystem 900, SNMP is not enabled on the storage system. IBM Spectrum® Control communicates with the SNMP agent on port 161 (UDP) to collect performance data.
To enable the SNMP agent, open the management GUI of a storage system, go to **Settings** > **Notifications** > **SNMP**, click **Agent**, and enter a community name.

The Data server is shut down automatically

If IBM Spectrum® Control cannot reach the database repository after 30 minutes, the Data server is shut down automatically.

Problem

You cannot complete actions in the IBM Spectrum Control GUI. This problem might occur when the Data server is down. The following message might be displayed:

BPCUI0209E A database operation cannot be completed.

Action

Verify that the local area network is available. Ensure that the IBM Spectrum Control database is up and running. Ensure that the Data server is up and running. Try the action again. For information about how to start DB2®, see [Manually starting Db2 on Windows](#). For information about how to start the Data server, see [Starting the IBM Spectrum Control servers by using scripts](#).

Fabric probe of DCFM CIM/OMs returns Java "Out of Memory" errors

During a fabric probe of a DCFM CIM/OM, if Java™ "Out of Memory" errors are returned, use the following steps to increase the server memory allocation for the DCFM CIM/OM.

Action

1. Open the DCFM management application. **Server** > **Options** > **Software Configuration** > **Memory Allocation**
2. Increase the server memory allocation to 1024.
3. Restart the DCFM services using the Server Management Console.
4. Run the fabric probe again.

Cluster resource group alerts are not triggered

This problem can occur on systems running PowerHA® SystemMirror® for AIX® or Microsoft Cluster Server (MSCS).

Problem

A computer alert is defined that uses one of the following triggering conditions:

- Cluster Resource Group Added
- Cluster Resource Group Removed
- Cluster Resource Group Moved

The event occurs, but no alert is triggered.

Action

The IBM Spectrum® Control agent checks to find out whether it is running in a IBM® PowerHA SystemMirror for AIX or MSCS cluster at the following times:

- When the agent is started
- When a probe is run

If the PowerHA SystemMirror for AIX or MSCS service is started after the agent, the agent is not aware of the cluster. Any cluster resource group alerts that you might have defined cannot be triggered.

To solve this problem, run a probe on one of the cluster nodes.

Slow performance when moving data from the database repository

Repocopy and other DB2® movement utilities might perform slowly or be in a wait state when moving data from the IBM Spectrum® Control database repository.

Problem

The IBM Spectrum Control user interface, repocopy tool, or DB2 data movement utilities might be in a wait state or perform poorly when you attempt to move data from the database repository while running discovery, probe, scan, and performance monitor jobs at the same time.

Action

Use the repocopy tool or DB2 data movement utilities to migrate data in the database repository during times when discovery, probe, scan, and performance monitoring jobs are not running or are scheduled to run.

Error message when running repocopy on server that uses remote database

This error occurs when you run **repocopy** on a server machine that uses a remote database.

Problem

You see these error messages when you run **repocopy** on a server machine that uses a remote database:

```
This application has failed to start because DB2APP.dll was not found.  
Re-installing the application may fix this problem.
```

```
Import/Export: Can't find library TSRMinsudb (TSRMinsudb.dll)  
in <library_path> or java.library.path  
<library_path>  
<java_library_path>
```

Action

Click **OK** to both messages.

The Data Server service fails with a logon failure when restarted

This error occurs when using Window's Active Directory domain policy.

Problem

When you restart the machine, the TSRMsrv1 user ID is no longer in the administrator's group and cannot log in to the machine. If there is a domain policy established that defined which users would be in the "Administrators" group, local users are automatically removed from the "Administrators" group when the system is restarted if the user ID is not one of the named accounts.

Action

Make sure that the user ID is one of the named accounts in the domain policy.

Specifying a LUN ID for the assignvol command

This error occurs when specifying a LUN ID for the **assignvol** command.

Problem

When specifying a LUN ID for the **assignvol** command, the LUN ID value will not be taken into consideration for the DS8000® storage devices. The reason for this is because the CIMOM does not offer support for client selectable LUN IDs.

Action

Understand the problem as described here.

Cannot monitor an EMC CIMOM on a Solaris server

This condition occurs when you try to monitor a Pegasus-based CIMOM like the EMC CIMOM.

Problem

When trying to log into a Pegasus-based CIMOM like the EMC CIMOM on a Solaris server, a login error is returned. If you are unable to log in, this could be a result of having the operating system language locale improperly set or not set at all. Check to see that the operating system language locale is properly set.

Action

An example of a properly set language locale for a Solaris server is found in the `/etc/default/init` file. This file contains the following lines which are dependant on the geographical location.

```
TZ=US/Arizona
CMASK=022
LC_COLLATE=en_US.ISO8859-1
LC_CTYPE=en_US.ISO8859-1
LC_MESSAGES=C
LC_MONETARY=en_US.ISO8859-1
LC_NUMERIC=en_US.ISO8859-1
LC_TIME=en_US.ISO8859-1
```

Incorrect entries are LC lines with incorrect locale entries or the LC language locale lines are missing.

SQLCODE-440 displayed if install IBM Spectrum Control on system with bad system clock

This error occurs when you install IBM Spectrum® Control on a system with a clock that is set for a future time and not the current time.

Problem

You receive this DB2® error message when you try to open the topology view for the fabric:

```
com.ibm.db2.jcc.c.SqlException: DB2 SQL error: SQLCODE: -440,
SQLSTATE: 42884, SQLERRMC: RTRIM;FUNCTION
```

Action

Correct the system clock and restart the IBM Spectrum Control services.

CIM agent runs slowly

This error occurs for the CIM agent.

Problem

The CIM agent runs slowly.

Action

Try stopping and restarting the IBM Spectrum® Control services. If this does not help, increase the value of `com.ibm.tpc.perf.ConteConnectTimeOut` in the following file:

```
installation_dir\device\conf\pm.conf
```

Assigned LUN is not recognized by the host

A LUN assigned to a host using IBM Spectrum® Control is not recognized by the host, and IBM Spectrum Control does not report the error.

Problem

This problem can occur if the host connection for the selected host was incorrectly configured on the storage subsystem, and the host is unable to log in to the storage subsystem. Some examples of configuration settings that, if incorrectly configured, can cause this problem are the storage subsystem I/O port configuration (FICON®/FcSf) or the OS selection.

Action

Remove all LUNs from the storage subsystem for the selected host and delete the host connection. Recreate the host connection with the correct configuration settings, and once again assign the LUNs to the host.

On AIX systems, numbers in the Storage Resource agent registry are printed as formatted numbers

On AIX® systems, numbers in the Storage Resource agent or configuration file are printed as formatted numbers.

Problem

Numbers are printed in formatted format in the following Storage Resource agent registry or configuration files:

```
/etc/Tivoli/TSRM/registryNA
<SRA_install_directory>/agent/config/Agent.config
```

For example, instead of 9510, the number appears as 9,510. This is a known in AIX libraries. This problem occurs with a earlier level C++ library (a level earlier than 7.0.0.5) on the system.

Action

Update your C++ libraries. To check your library version, run the following command:

```
lsldpp -L | grep -i xlc
```

An example of the output is shown:

xlc.aix50.rte	7.0.0.3	C	F	C Set ++ Runtime for AIX 5.0
xlc.cpp	6.0.0.0	C	F	C for AIX Preprocessor
xlc.msg.en_US.cpp	6.0.0.0	C	F	C for AIX Preprocessor
xlc.msg.en_US.rte	7.0.0.0	C	F	C Set ++ Runtime
xlc.rte	7.0.0.1	C	F	C Set ++ Runtime

To upgrade your libraries, follow these steps:

1. Get the library updates. Go to <http://www-933.ibm.com/support/fixcentral/>.
2. Select the following:

Product Group	Rational
Product	XL C++ Runtime
Installed Version	8.0.0.0
Platform	AIX

3. Click **Continue**, then click **Continue** again.
4. Download the updates.

Fabric probe job failed

The fabric probe job failed.

Problem

A fabric probe job failed. Here is an example of the error messages you see in the fabric probe log:

```
3/13/09 12:53:32 PM BTADS0029I Scanner AttributePEOnly data from agent
x.xx.xxx.xxx:9510 has not changed since last scan.
3/13/09 12:53:33 PM BTADS0033E Error invoking AttributePEOnly on host
x.xx.xxx.xxx:9510 .
3/13/09 12:53:33 PM java.lang.reflect.UndeclaredThrowableException
at $Proxy16.invoke(Unknown Source)
at com.ibm.tpc.discovery.tsanm.InbandScanner.process(InbandScanner.java:136)
at com.ibm.tpc.infrastructure.threads.TPCThread.run(TPCThread.java:257)
Caused by: java.net.ConnectException: Connection refused: connect
at java.net.PlainSocketImpl.socketConnect(Native Method)
...
3/13/09 12:53:46 PM BTADS0033E Error invoking Topology on host
x.xx.xxx.xxx:9510 .
3/13/09 12:53:46 PM java.lang.reflect.UndeclaredThrowableException
at $Proxy16.invoke(Unknown Source)
at com.ibm.tpc.discovery.tsanm.InbandScanner.process(InbandScanner.java:136)
at com.ibm.tpc.infrastructure.threads.TPCThread.run(TPCThread.java:257)
```

```
Caused by: java.net.SocketException: Operation timed out: connect:could be
    due to invalid address
    at java.net.PlainSocketImpl.socketConnect(Native Method)
```

Action

Check the IP address in the /etc/hosts file on the agent system. If the IP address has changed, enter the correct IP address and restart the agent on the system.

Error Messages: Get SQLCODE: -964, SQLSTATE: 57011, Unexpected error occurred while performing [a function]

You might receive these error messages when performing tasks within IBM Spectrum® Control.

Problem

When performing tasks within IBM Spectrum Control that require accessing and storing data in the database repository, you might receive error messages similar to the following:

```
[message #]: Unexpected error occurred while performing the following [functions]
```

where [message

#] represents the number of the error message and [functions] represents a task that you were performing.

You might also see the following error message in the database installation log:

```
Instruction SQL : insert into t_stat_file_temp
SQLSTATE: 57011, Vendor error code: -964
DB2 SQL error: SQLCODE: -964, SQLSTATE: 57011,
SQLERRMC: null
```

These error messages are a result of running out of space in the transaction log. The transaction log configuration is set as:

```
db2 update db cfg for $DBNAME using logprimary 8
db2 update db cfg for $DBNAME using logsecond 16
```

Each transaction log file is 10 MB (logFileSize=2500 pages of 4 KB each). Eight (8) primary log files are allocated all the time; secondary log files are allocated as needed, and deleted when not needed, to the maximum of 16.

Action

Increase the transaction log files. See the DB2® documentation for information about increasing the transaction log files.

Data server crashes on AIX

This error occurs when you are running the Data server on AIX®.

Problem

The Data server crashes on AIX with the following messages in the system error logs:

```
ERROR LOGGING BUFFER OVERFLOW caused by
EXCESSIVE LOGGING BY SOFTWARE PROGRAM.
```

Action

Increase the paging space on the server.

No performance data is retrieved

This error occurs when you try to get performance data.

Problem

You are not successful in getting performance data.

Action

For a successful performance data collection to occur, the device must have been discovered and probed successfully.

Performance monitoring job fails immediately after it starts

This error occurs when you run a performance monitoring job and it fails immediately after it starts.

Problem

The performance monitoring job fails immediately after it starts.

Action

This problem usually occurs when there is a network connectivity problem with the CIM agent.

Error message "performance data files could not be correlated"

This error occurs when you uninstall the agent using the IBM Spectrum® Control uninstallation program.

Problem

You get this error message for the SAN Volume Controller subsystem during performance data collection.

Action

For SAN Volume Controller subsystems, this error message can occur during performance data collection. This is usually a problem with mismatching timestamps for the I/O statistics dump files on the SAN Volume Controller itself. This problem can be resolved by setting the time zone properly on the cluster.

Performance correlation step takes a long time

This error occurs when a performance correlation step is done.

Problem

The performance correlation step takes a long time.

Action

The performance correlation step can take a long time to complete if there are a lot of volumes. This causes a delay in collecting the first performance sample. To work around this problem, place the CIM agent and IBM Spectrum® Control servers on faster machines. You can also use multiple CIM agents to monitor a set of storage subsystems.

Cannot define a performance monitor for a switch

This error occurs when trying to define a performance monitor for a switch.

Problem

You cannot define a performance monitor for a switch.

Action

If you cannot define a performance monitor for a switch because it does not appear in the list of switches that can be monitored, do the following:

- Check to see that the SMI-S version supported by the switch vendor is 1.1 or higher.
- Check to see if the CIM agent supports the switch sub-profile.
- Check to see if the CIM agent discovery completed successfully.
- Check to see if a fabric probe was run on the fabric that contains the switch after the CIM agent discovery completed successfully for that switch. Then verify that this fabric probe also completed successfully.

File system storing database tables runs out of space

This problem occurs on systems running AIX® 5.x.

Problem

The file system storing the database table spaces runs out of space. This is shown by the following kinds of errors as seen in the DB2® log called `db2diag.log`:

```
2005-11-06-12.59.10.815891-420 E36562387C690      LEVEL: Error
PID       : 176316          TID    : 1          PROC  : db2pclnr 0
INSTANCE: db2inst1        NODE   : 000
FUNCTION: DB2 UDB, buffer pool services, sqlbClnrAsyncWriteCompletion,
probe:0
MESSAGE : ADM6017E The table space "TPCTBSPTMP" (ID "5") is full.
          Detected on
          container "/home/db2inst1/db2inst1/TPCDB/TPC/TPCTBSPTMP"
          (ID "0").
          The underlying file system is full or the maximum allowed
          space usage
          for the file system has been reached. It is also possible
          that there
          are user limits in place with respect to maximum file size
          and these limits have been reached.
```

This problem is seen on AIX 5.x with JFS file system. Large files need to be enabled on the file system because the table space files can get quite large. Enable large files on JFS or JFS2.

Action

During installation of AIX 5.1 or later, ensure that JFS or JFS2 with large file system support is installed. Follow the AIX documentation for this purpose as well as to back up and restore data between file systems. It is recommended to ensure that the file system size (especially for `/home`) is not too small (for example, not less than 4 GB). Back up your existing data before you recreate your file system. For information about file systems, see: http://www-01.ibm.com/support/knowledgecenter/ssw_aix_53/com.ibm.aix.baseadm/doc/baseadmndita/fs_types.htm.

Alternatively, you can set the `DIAGPATH` and `SPM_LOG_PATH` variables to point to an external, fast disk drive.

Error message: There is a problem with this website's security certificate

You receive this error message on Internet Explorer when you launch the DS8000® Element Manager.

Problem

You cannot launch the DS8000 Element Manager.

Action

To resolve this issue, follow these steps:

1. Launch the DS8000 Element Manager in a Web browser.
2. You see this error message:

There is a problem with this website's security certificate.

3. Click the link labeled **Continue to this website (not recommended)**.
Note: In the Internet Explorer taskbar, you see a red box labeled **Certificate Error**. Click the **Certificate Error** box.
4. Click **View certificates** in the dialog window. You see a **Certificate** dialog window.
5. Click **Install Certificate** button.
6. Click **Next**.
7. Select **Place all certificates in the following store** radio button.
8. Click **Browse**. You see the **Select Certificate Store** dialog window.
9. Select the **Show physical stores** check box.
10. Expand the **Trusted Root Certification Authorities** navigation tree entry.
11. Select the **Local Computer** subtree entry.
12. Click **OK**.
13. Click **Next**.
14. Click **Finish**.
15. Click **OK** on the message box **The import was successful**.
16. Click **OK** to close the **Certificate** window.
17. Close the Web browser.
18. Launch the DS8000 Element Manager.

Cannot discover fabric on Red Hat Linux using Emulex HBA

A fabric probe operation performed on Red Hat® Linux® using a Storage Resource agent cannot find the Emulex HBA.

Problem

When you installed the Emulex driver, the installation might not have installed the 32-bit DFC libraries because of missing prerequisites. The following messages appear during the Emulex driver installation:

```
Installing DFC libraries via RPM ...
Unable to install 32-bit development kit library due to missing
32-bit shared object library: libnl.so
If 32-bit development kit desired, please install RPM package:
libnl-1.0-0.10.pre5.4.i386 (or equivalent) that includes 32-bit libnl.so.
```

Action

Verify that the DFC 32-bit library, `/usr/lib/libdfc.so`, is on the system. If the library is missing, install the missing prerequisite package shown during the Emulex driver installation, and rerun the Emulex driver installation.

Cannot get version of Storage Resource agent on SUSE Linux

You cannot get the version of the Storage Resource agent on SUSE Linux®.

Problem

You try to run the following command to get the version of the Storage Resource agent on SUSE Linux:

```
./agent.sh version
```

This command does not get the version of the Storage Resource agent.

Action

To get the version of the Storage Resource agent, run the following command:

```
./agent.sh versionall
```

Cannot log on to two instances of IBM Spectrum Control by using the same browser

You cannot log on to two instances of IBM Spectrum® Control by using the same browser.

Problem

When you try to log on to more than one instance of IBM Spectrum Control by adding a tab or opening a new window in a browser, you are logged out of the first instance. This situation occurs because the Lightweight Third-Party Authentication (LTPA) cookie that is set by the first server (`TPC_server1`) is being overwritten by the LTPA cookie that is set by the second server (`TPC_server2`).

Action

To avoid this issue, take one of the following actions:

- Use different web browsers to log on to the different servers. For example, you can log on to one instance on Internet Explorer and the second instance by using Mozilla Firefox.
- If you are using single sign-on, make sure the Jazz® for Service Management servers are configured with the same LDAP information. By configuring the second authentication service to use the same LTPA keys as the first server, the LTPA cookies do not overwrite each other.
- If you are not using single sign-on, configure the LTPA cookies on each server to use the fully qualified host name of the server instead of only the parent domain name. Complete the following steps for each server to configure the domain name:

1. Log on to the IBM Spectrum Control server with administrative privileges.
2. Open a command prompt and go to `installation_dir/wlp/usr/servers/webServer` directory.
3. Edit the `server.xml` file and change this line:

```
<webAppSecurity ssoUseDomainFromURL="true" />
```

to:

```
<webAppSecurity ssoDomainNames="<fully qualified host name of
IBM Spectrum Control server>" />
```

4. Save and exit the `server.xml` file.

Because the servers are now identified by their fully qualified host names, an LTPA cookie that is set by one server does not overwrite an LTPA cookie that is set by the other server.

Warning in Probe Log for Storage Resource agents on AIX (STA0115W)

If certain services, such as `mountd` daemon, are not running on the AIX® Storage Resource agents, warning entries are displayed in the Probe log for Storage Resource agents deployed on AIX.

About this task

To review the Storage Resource agent Probe log for warnings and fix the underlying issue, complete the following steps:

Procedure

1. Run the **'rpcinfo -p localhost'** command to ensure that the **portmapper**, **status**, **nlocmgr**, **nfs**, and **mountd** programs are registered with the portmapper. The **'rpcinfo -p localhost'** command output shows that the **portmapper**, **status**, **nlocmgr**, **nfs**, and **mountd** programs are registered with the portmapper and up and running.
2. If you do not see the programs, start them.
For example, to start **mountd** on the AIX operating system, run the following command:

```
'startsrc -s mountd'
```

3. Run the **'rpcinfo -p localhost'** command again to verify that the required daemons are running.

Example

Here is an example of a message that appears in the log:

```
Apr 4, 12 1:44:11.000 AM EST AGT0446I Fabric discovery found switch fabric
Apr 4, 12 1:44:12.000 AM EST STA0115W Cannot connect to mountd to list exports:
RPC: 1832-019 Program not registered
Apr 4, 12 1:44:12.000 AM EST STA0110I Probe completed with warnings
```

A generic Storage Resource agent probe error might be a registration problem / GUID mismatch

A generic Storage Resource agent (SRA) probe error might occur when sending data to the server.

Problem

A generic Storage Resource agent (SRA) probe error might occur when sending data to the server, with the following in the SRA probe log:

```
STA0109I Probe completed successfully
STA0249I Sending results to server
AGT0306E Failed to send Probe complete status to server, for Agent in /opt/IBM/TPC
STA0251E Results not saved
```

Action

To check if this problem is caused by a registration problem or GUID mismatch, review the Data Server logs (TPC\data\log\TPCD_0000xx.log file) for the following:

```
SRV0040E: No agent is registered under machine identifier
XX.XX.X.XX (GUID: XXXxxXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX)

GEN0100E: Error processing request from host
<XX.XX.X.XX >, for service <Scheduler>, request(3, 1)
```

To resolve this issue manually register the SRA again so the GUID matches. Register the SRA by running the following command on the SRA host:

```
<SRA_install_location>/agent/bin/Agent -notify register -serverName <server_name>
```

Following IBM Spectrum Control server upgrade, the Storage Resource agents (SRA) might lack the Upgrade action

Immediately following an upgrade of the IBM Spectrum® Control server, the SRAs will have the Agent State of *Upgrade needed*, but might not have the option to start the SRA upgrade.

Problem

An Agent State of *Upgrade needed* occurs following an upgrade of the IBM Spectrum Control server, but there is no option available in IBM Spectrum Control to start the SRA upgrade.

Action

When a probe is run on the SRA, the option to upgrade will be available. Simply wait for the scheduled probes to occur, or to speed up the process, manually start probes of the SRAs, and then upgrade each to the current version.

Related tasks

- [Upgrading Storage Resource agents](#)

How to fix IBM Storwize V7000 Unified authentication errors

If the IP address for IBM® Storwize® V7000 Unified is not configured correctly on SAN Volume Controller, you get an authentication error when you add the storage system in IBM Spectrum® Control.

Problem

You can add IBM Storwize V7000 Unified on a SAN Volume Controller to IBM Spectrum Control as either a file storage system, or block storage system, or both. However, if you add a block storage system IP address for a file storage system, or vice versa, an authentication error occurs.

Action

To solve the issue, complete these steps:

1. Get the IP address of IBM Storwize V7000 Unified storage system from the Storwize V7000 Unified Management GUI.
2. Get the IP address for the Storwize V7000 Unified storage system registered on SAN Volume Controller by using the following command on a command line of SAN Volume Controller. **svcinfo lscluster -devicename**
3. Update the IP address on SAN Volume Controller with the IP address from step 1 by using the following command: **chsystemip -clusterip <IP address>**
4. Retry the connection on IBM Spectrum Control.

Slow performance on VMware virtual machines due to shared VMDK

The VMware virtual machine where you run IBM Spectrum® Control exhibits slow performance.

Problem

The VMware virtual machine where you run IBM Spectrum Control is using a shared VMDK (Virtual Machine Disk) with other virtual machines. When you use a shared VMDK, it reduces the total number of disk IOPS (Input/Output Operations Per Second) available for each virtual machine; this causes poor performance and high disk latency on the virtual machine where you run IBM Spectrum Control.

Action

Use a dedicated VMDK for the VMware virtual machine where you run IBM Spectrum Control. Reconfigure the other virtual machines to *not* use the same VMDK as your IBM Spectrum Control virtual machine.

For more information, see <https://kb.vmware.com/s/article/1031773>.

How to disable cipher block chaining (CBC) ciphers in IBM Spectrum Control in response to the Lucky 13 vulnerability.

If needed, you can disable cipher block chaining (CBC) ciphers in IBM Spectrum® Control in response to the Lucky 13 vulnerability.

Problem

IBM Spectrum Control contains all of the necessary patches to properly address the Lucky 13 vulnerability (<https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2013-0169>). However, you might also want to disable CBC ciphers in IBM Spectrum Control in response to the Lucky 13 vulnerability.

Important: After you disable the CBC ciphers, IBM Spectrum Control does not manage IBM® DS8000® storage systems unless they are updated to a firmware level that provides GCM cipher support. For more information, contact IBM DS8000 storage system support.

Action

To disable CBC ciphers in IBM Spectrum Control, complete these steps:

1. Stop all IBM Spectrum Control servers.
2. Open the `installation_dir/jre/lib/security/java.security` file.
3. To disable CBC ciphers in IBM Spectrum Control, insert `AES_128_CBC, AES_256_CBC` to the `jdk.tls.disabledAlgorithms` line.

BEFORE

```
jdk.tls.disabledAlgorithms=MD5withRSA, DH keySize < 1024, EC keySize < 224, TLSv1, TLSv1.1, anon, NULL
```

AFTER

```
jdk.tls.disabledAlgorithms=MD5withRSA, DH keySize < 1024, EC keySize < 224, TLSv1, TLSv1.1, AES_128_CBC, AES_256_CBC, anon, NULL
```

4. Restart the IBM Spectrum Control servers.

Related information

- [fqz0 t_starting_tpc_services.html](#)

Db2 and database troubleshooting

Use this section to troubleshoot and resolve problems with IBM® Db2® and the database.

- [Db2 crashes on RHEL 7.2](#)
IBM Db2 servers might crash when the RemoveIPC option for Red Hat® Enterprise Linux® 7.2 is set to yes.
- [Db2 log messages ADM5530W can be ignored](#)
ADM5530W messages in the Db2 Event Log and db2diag.log indicate that the materialized query table (MQT) is using the 'NOT LOGGED INITIALLY' option to improve database performance and create fewer transaction log entries. These messages do not require any user action and can be ignored.
- [Db2 log files fill the C drive on the Windows DB2 system](#)
On Db2 servers that run for extended periods of time, the Db2 logs can grow large enough to fill the C: drive if the default Db2 log location is used.
- [Cannot read Korean language logs from GUI](#)
This problem occurs in Korean language environments.
- [Installation fails with "No valid local database found" error](#)
During the IBM Spectrum® Control installation, an error occurs stating that no valid databases were found.
- [Error message: GUI0023M](#)
This error occurs when performing a system reboot.
- [Error message: SRV0024E](#)
This error occurs when Db2 is down or not running.
- [Error message: SRV0044E](#)
This error occurs when the password for the Db2 user ID account for which IBM Spectrum Control is installed under is changed.
- [Db2 instance crashes](#)
This problem occurs when running an agent probe.
- [Db2 connection errors](#)
These errors occur when rebooting remote Db2 servers.
- [Error message: SQL0968C](#)
This error occurs when the file system storing the database table runs out of space.
- [Error message when installing Db2 on Windows](#)
You get an error message when installing Db2 on a Windows system.
- [Error message: HWNOP0033 Database operation failed](#)
You receive this error message after running several storage optimizer jobs.
- [Error message: DB2 SQL error: SQLCODE: -964, SQLSTATE: 57011](#)
You receive this SQL exception when you migrate IBM Spectrum Control data.
- [Error message: DB2 SQL error: SQLCODE: -973, STATE: 57011, SQLERRMC: PCKCACHESZ](#)
You receive this error when opening a panel in the IBM Spectrum Control user interface that requires information from the database repository. For example, when you generate a report or view a list of devices managed by CIMOMs.
- [Error message: DEBUG: Error 2836](#)
After you install Db2, you might see this error message in the Db2 installation log file.
- [Error message: SRV0046E, Db2 stops running](#)
You receive this error message when Db2 stops running.
- [Error message: The statement is too long or too complex \(SQL0101N\)](#)
This error occurs when there is not enough storage available in the database heap to process a statement.
- [Db2 expired license error](#)
IBM Spectrum Control installation and operation requires that Db2 is running. After the Db2 trial license period expires, Db2 does not start without a valid license.
- [Context root not found on IBM Spectrum Control 5.3.5 or later on Red Hat Enterprise Linux 7](#)
On the Red Hat Enterprise Linux 7 operating system, the Db2 and IBM Spectrum Control web server are configured to start at the same time upon start up.

Db2 crashes on RHEL 7.2

IBM® Db2® servers might crash when the RemoveIPC option for Red Hat® Enterprise Linux® 7.2 is set to yes.

Problem

The logind.conf file for RHEL contains the option **RemoveIPC**. That option controls whether System V and POSIX IPC objects that belong to the user are removed when the user logs out. For RHEL 7.2, **RemoveIPC** is set to yes by default, which can cause the Db2 database to crash every time a database administrative task cron session finishes.

When Db2 crashes, you might see the following message in the db2diag.log file:

```
MESSAGE: Unexpected OS error. This most likely means that resources have been  
torn down from underneath the prefetcher. Terminating the prefetcher  
to prevent infinite looping.
```

Action

To resolve this problem, complete the following steps:

1. On the computer where RHEL 7.2 and Db2 is installed, open /etc/systemd/logind.conf in a text editor.
2. Set **RemoveIPC** to **no**.
3. Save the file.

You might need to restart RHEL 7.2 for the change to take effect.

For more information about this problem, see [Applications using IPC \(semaphores, shared memory, message queues\) have problems after update to RHEL 7.2.](#)

Db2 log messages ADM5530W can be ignored

ADM5530W messages in the Db2® Event Log and db2diag.log indicate that the materialized query table (MQT) is using the 'NOT LOGGED INITIALLY' option to improve database performance and create fewer transaction log entries. These messages do not require any user action and can be ignored.

To prevent these messages from appearing in the Event Log and db2diag.log, set the Db2 [`notifylevel`] and [`diaglevel`] parameters to 2.

Related information

- 🔗 [notifylevel - Notify level configuration parameter](#)
- 🔗 [diaglevel - Diagnostic error capture level configuration parameter](#)

Db2 log files fill the C drive on the Windows DB2® system

On Db2® servers that run for extended periods of time, the Db2 logs can grow large enough to fill the C: drive if the default Db2 log location is used.

Problem

When installing Db2 on a Microsoft Windows system, the default location of the DB2DIAG and SPMLOG log files is the C: drive. After running IBM Spectrum® Control for an extended period of time, the Db2 logs can fill up the C: drive, making it difficult to install other applications on that system. This can also prevent other applications from writing to the C: drive.

Action

Use the information provided in the following links to manage your Db2 logs:

- Control the size of the Db2 diagnostic log using the **diagsize** database configuration parameter: [diagsize - Rotating diagnostic and administration notification logs configuration parameter](#).
- Archive the Db2 diagnostic logs using the **archive** database configuration parameter: [db2diag - db2diag logs analysis tool command](#).

Cannot read Korean language logs from GUI

This problem occurs in Korean language environments.

Problem

You will not be able to view the log files from the IBM Spectrum® Control GUI because of a Db2® JDBC driver problem.

Action

You will have to read the files using another text viewer like **Notepad**.

Installation fails with "No valid local database found" error

During the IBM Spectrum® Control installation, an error occurs stating that no valid databases were found.

Action

This problem can occur if the db2profile is not sourced correctly, or if Db2® has not installed successfully. To resolve this issue, perform the following steps:

1. If the installation is on a UNIX server, be sure to source the db2profile. For example, for an instance of **db2inst1**, source the db2profile by entering the following: `. /home/db2inst1/sqllib/db2profile`. Note the space between the `.` and `/home`.
2. Verify the Db2 installation. See [Verifying that Db2 is installed correctly](#) for information about how to verify a Db2 installation.

Error message: GUI0023M

This error occurs when performing a system reboot.

Problem

After a reboot of the system, you receive the following error message:

```
GUI0023M Unable to connect to server after reboot.
```

If abnormal behaviors are observed which might be related to or impact Db2®, be sure to stop your Db2 instance properly before you perform a system reboot. If you do not stop Db2 properly, Db2 might take a long time to start. This will also cause IBM Spectrum® Control to not start.

Action

The Db2 instance might not have stopped properly. Manually stop and restart Db2.

Error message: SRV0024E

This error occurs when Db2® is down or not running.

Problem

This might be a problem with Db2. The Db2 database can be down or Db2 is not running. Check to see if Db2 is running. You will see this message:

```
PM SRV0024E The requested service provider has been shutdown.
```

Action

If Db2 is not running, start Db2 and then stop and restart the Device server and Data server.

Error message: SRV0044E

This error occurs when the password for the Db2® user ID account for which IBM Spectrum® Control is installed under is changed.

Problem

The GUI will not start and an error occurs in the log:

```
PM SRV0044E: Unable to connect to repository database. Please ensure that the repository database is up and running. Connection authorization failure occurred. Reason: password invalid.
```

Action

To change the Db2 user ID password, follow these steps:

1. From a command prompt, change to the directory where the file **repository.config** resides in the IBM Spectrum Control installation path. For example on Windows:

```
<installation_dir>\data\config
```

2. Run the following command for the Data server:

```
java -classpath "<installation_dir>\data\server\lib\TSRMsrv.zip" com.tivoli.itsrm.repository.Transform -p <new_db2userid_password> repository.config
```

3. Make sure the Device server is running. Run the following commands for the Device server:

```
cd <installation_dir>\device\bin\<platform> setenv srmcp -u <user_ID> -p <password> ConfigService setPw <new_password>
```

Where **<new_password>** is the new password. This sets the Db2 user ID password (not the Db2 administrator password).

4. Stop and then restart IBM Spectrum Control - Data server from the Windows Services window.
5. Restart IBM® WebSphere® Application Server V6 - Device server from the Windows Services window.
6. Start the IBM Spectrum Control GUI.

Db2 instance crashes

This problem occurs when running an agent probe.

Problem

An agent probe causes the Db2® instance to crash. You also receive this message in the Db2 logs:

```
ADM0501C A stack overflow exception has occurred. The DB2 instance has terminated abnormally.
```

Action

To resolve this problem, increase the Db2 stack size. Go to the Db2 directory:

C:\Program Files\IBM\SQLLIB

Follow these steps:

1. Stop the Device server and Data server.
2. Stop the Db2 database.
3. Run the following command:

```
.\misc\db2hdr .\bin\db2syscs.exe /s: 512
```

4. Start the Db2 database.
5. Start the Device server and Data Server.

Db2 connection errors

These errors occur when rebooting remote Db2® servers.

Problem

Getting Db2 connection errors in IBM Spectrum® Control server log files

Action

This could be caused by the remote Db2 server being rebooted. If this is the case, restart the Data Server and Device server.

Error message: SQL0968C

This error occurs when the file system storing the database table runs out of space.

Problem

You see this error message from Db2®:

SQL0968C The file system is full.

This is the explanation from Db2:

One of the file systems containing the database is full. This file system may contain the database directory, the database log files, or a table space container.

The statement cannot be processed.

Action

Free system space by erasing unwanted files. Do not erase database files. If additional space is required, it may be necessary to drop tables and indexes identified as not required.

On UNIX-based systems, this disk full condition may be due to exceeding the maximum file size allowed for the current user ID. Use the **chuser** command to update the **fsize**. A reboot might be necessary.

This disk full condition may be caused when containers are of varying sizes. If there is sufficient space in the file system, drop the table space and recreate it with containers of equal size.

For AIX® systems, check to see if the AIX file size limit goes beyond 2 GB. Check the sizes of the files. As the root user, run this command:

```
lsfs -qa
```

If you see "bf: false", then you need to enable the large file option to allow files to go over 2 GB.

It is suggested that you monitor Db2's database performance. You can do this through DB2®'s Health Center. For more information about DB2's Health Center, see [Planning to monitor Db2](#) and [IBM® Db2 Version 11.5 Knowledge documentation](#). Click **Monitoring > Database systems > Monitoring database systems > Monitoring database health**.

Error message when installing Db2 on Windows

You get an error message when installing Db2® on a Windows system.

You can receive one of the following error messages:

An error occurred while updating the services file on the system for the service name 'db2c_DB2' with port '25000'.

No valid local database found on the system.

You receive this error message because Db2 failed to configure the port at installation time.

An error occurred while initializing the Profile Registry.

Error message: HWNOP0033 Database operation failed

You receive this error message after running several storage optimizer jobs.

Problem

You receive this error message after running several storage optimizer jobs:

```
HWNOP0013E Database operation failed com.ibm.db2.jcc.b.SqlException:  
DB2 SQL error: SQLCODE: -956, SQLSTATE: 57011, SQLERRMC: null
```

This message indicates that all available memory for the database has been used. There is not enough storage available in the database heap to process the statement.

Action

Increase the database heap size to allow a larger database heap. If the database heap size is set to AUTOMATIC, you need to increase either the DATABASE_MEMORY database configuration setting, or the INSTANCE_MEMORY database manager configuration setting.

When updating the configuration parameters, it is suggested to change them by 10% of the current size until the error condition is resolved.

Here is an example of how to change the database heap size to 2400:

```
db2 UPDATE DB CFG FOR TPCDB USING DBHEAP 2400  
(the IBM Spectrum  
Control database schema sets the  
database heap size to 1800)
```

Here is an example of how to change the database heap size when disconnected from the database:

```
db2 CONNECT RESET;  
db2 UPDATE DB CFG FOR TPCDB USING DBHEAP 2400;
```

If all associated configuration parameters are set to either AUTOMATIC or COMPUTED, and the memory demands of the instance exceed the amount of memory configured on the system, then possible solutions would be to reduce the database workload or adding additional memory to the system.

Error message: DB2® SQL error: SQLCODE: -964, SQLSTATE: 57011

You receive this SQL exception when you migrate IBM Spectrum® Control data.

Problem

When migrating IBM Spectrum Control data, you get this error message in the database installation log:

```
Instruction SQL : insert into t_stat_file_temp  
SQLSTATE: 57011, Vendor error code: -964  
DB2 SQL error: SQLCODE: -964, SQLSTATE: 57011,  
SQLERRMC: null
```

This message indicates that there is not enough storage available in the database heap to process the statement. This error message is a result of running out of space in the transaction log. The transaction log configuration is set as:

```
db2 update db cfg for $DBNAME using logprimary 8  
db2 update db cfg for $DBNAME using logsecond 100
```

Each transaction log file is 10 MB (logFileSize=2500 pages of 4 KB each). Eight primary log files are allocated all the time. Secondary log files are allocated as needed, and deleted when not needed, to the maximum of 100.

Action

Increase the transaction log files. See the Db2® documentation for information about increasing the transaction log files. Go to [IBM® Db2 Version 11.5 Knowledge documentation](#). Search for **logprimary** and **logsecond**.

Error message: DB2® SQL error: SQLCODE: -973,STATE:57011, SQLERRMC:PCKCACHESZ

You receive this error when opening a panel in the IBM Spectrum® Control user interface that requires information from the database repository. For example, when you generate a report or view a list of devices managed by CIMOMs.

Problem

When you open a panel that requires information from the database repository, you might receive the following error:

```
Server: server_name      Status: 0
SQLSTATE: 57011, Vendor error code: -973
DB2 SQL error: SQLCODE: -973, SQLSTATE: 57011,
SQLERRMC: PCKCACHESZ
```

This message indicates an error related to the PCKCACHESZ configuration setting in Db2®. This parameter is allocated out of the database shared memory, and is used for caching of sections for static and dynamic SQL and XQuery statements on a database.

Action

Increase the value for PCKCACHESZ until the error in the user interface is resolved. Initially, increase the value of PCKCACHESZ to 78525. If that does not resolve the error, increase the value in 10% increments until the error is resolved.

To increase the value for PCKCACHESZ, complete the following steps:

1. Check the current value for PCKCACHESZ by running the following commands at the Db2 prompt:

```
db2 connect to tpcdb
db2 get db cfg show detail | grep PCKCACHESZ
```

2. Increase PCKCACHESZ to 78525:

```
db2 update db cfg for TPCDB using PCKCACHESZ 78525
```

Note: If PCKCACHESZ is set to **AUTOMATIC**, all applications must disconnect from the Db2 in order for the change to take effect.

3. If increasing PCKCACHESZ to 78525 does not resolve the problem, increase the value by 10% increments until the error condition is resolved:

```
db2 update db cfg for TPCDB using PCKCACHESZ NEW_VALUE
```

where the initial **NEW_VALUE** is 10% greater than 78525.

Error message: DEBUG: Error 2836

After you install Db2®, you might see this error message in the Db2 installation log file.

Problem

The error message is:

```
DEBUG: Error 2836: The control image_noJava on the dialog SetupInitialization
can not take focus
Internal Error 2836. SetupInitialization, image_noJava
Action 14:29:48: SetupInitialization. Dialog created
Action ended 14:29:48: SetupInitialization. Return value 1.
```

Action

This is an informational message for the Windows installation program and will not cause problems.

Error message: SRV0046E, Db2 stops running

You receive this error message when Db2® stops running.

Problem

You receive this error message:

```
SRV0046E: Repository connection status check failed.
The database manager is not able to accept new requests, has terminated all
requests in progress, or has terminated this particular request due to
unexpected error conditions detected at the target system.
```

This issue can occur on IA32 32-bit platforms due to a timing problem.

Action

Restart Db2.

Error message: The statement is too long or too complex (SQL0101N)

This error occurs when there is not enough storage available in the database heap to process a statement.

Problem

When you add many columns to a custom report, the following error message might be displayed:

```
SQL0101N The statement is too long or too complex
```

This message indicates that all available memory for the database is used.

Action

Increase the database heap size to allow a larger database heap. For example, to change the IBM Spectrum® Control database heap size, use the **db2 UPDATE DB** command:

```
db2 UPDATE DB CFG FOR TPCDB USING STMHEAP 500000
```

Db2 expired license error

IBM Spectrum® Control installation and operation requires that Db2® is running. After the Db2 trial license period expires, Db2 does not start without a valid license.

Problem

You get an error when you attempt to install or operate IBM Spectrum Control because Db2 is not running. After the Db2 trial license period expires, Db2 requires a valid license in order to start.

Action

Retrieve the valid Db2 license that is supplied on the IBM Spectrum Control download site. Find the db2awse_o.lic file in the awse_o/db2/license/ directory and run the db2licm command to apply the new license. For example, from the command line in the license directory run the following command:

```
db2licm -a db2awse_o.lic
```

Related information

- 📄 [IBM Spectrum Control version downloads](#)

Context root not found on IBM Spectrum Control 5.3.5 or later on Red Hat Enterprise Linux 7

On the Red Hat® Enterprise Linux® 7 operating system, the Db2® and IBM Spectrum® Control web server are configured to start at the same time upon start up.

Problem

If Db2 takes significantly longer to start up than the IBM Spectrum Control web server, you see a Context Root Not Found page when trying to load the IBM Spectrum Control GUI.

Action

Restarting the web server after Db2 has fully started might resolve the issue until the next start up. If not, to fix the issue across start ups, complete the following steps:

1. Edit the /etc/systemd/system/webWLP.service directory by adding the lines in **bold**:

```
[Unit]
Description=IBM Spectrum Control Web Server
After=db2fmcd.service

[Service]
ExecStartPre=/usr/bin/sleep 45
ExecStart=/usr/bin/bash -c "/opt/IBM/TPC/wlp/bin/server start webServer --clean >
/opt/IBM/TPC/wlp/bin/../../web/log/startServer.log"
ExecStop=/usr/bin/bash -c "/opt/IBM/TPC/wlp/bin/server stop webServer >
/opt/IBM/TPC/wlp/bin/../../web/log/stopServer.log"
TimeoutStopSec=360
Type=forking
LimitNOFILE=4096
LimitNPROC=131072
LimitFSIZE=infinity
LimitCORE=infinity

[Install]
WantedBy=multi-user.target
```

2. Run the following command as root:

```
# systemctl daemon-reload
```

Ensure that you also implemented this fix for Db2 AutoStart on Red Hat Enterprise Linux 7. See [DB2® AutoStart is not supported on Red Hat Enterprise Linux version 7](#).

Installation, uninstallation, and upgrading

Use this section to troubleshoot and resolve problems with IBM Spectrum® Control installation, uninstallation, and upgrading.

- [Creating a keystore for an IBM Spectrum Control server](#)
You need this recovery procedure if you do not know the keystore password for the Alert server, Device server, or Web server during an upgrade of IBM Spectrum Control.
- [Error message: INS3105E Failure to install remote agent](#)
This error occurs when you are installing a remote agent on a UNIX or Linux® system. .
- [Installing a remote agent fails](#)
This error occurs when you install a remote agent and it fails.
- [The Solaris HBA identifies a SCSI device instead of fibre channel LUNs](#)
This problem occurs when the Oracle Solaris HBA identifies a SCSI device instead of fibre channel LUNs.
- [Installing IBM Spectrum Control on Windows with installation files on a Samba fileshare](#)
UNIX and Linux file systems use Samba software to mount file systems on Windows. UNIX and Linux mark hidden files with a leading dot '.' in the file or the folder name.
- [IBM Spectrum Control agent installation fails when files are mapped to a drive in a Windows Terminal Services environment](#)
The IBM Spectrum Control agent installation fails when files are mapped to a drive in a Windows Terminal Services environment.
- [Db2 services do not start when you restart the system on a UNIX platform](#)
The Db2® services do not start when you perform a system reboot on UNIX.
- [Cannot install the Storage Resource agent on AIX through RSH](#)
You cannot install the Storage Resource agent on an AIX® system that has the Remote Shell (RSH), but does not have the Secure Shell (SSH).
- [Storage Resource agent failed to install using install command](#)
The Storage Resource agent failed to install when using the install command.
- [Upgrading to IBM Spectrum Control 5.4.1 on RHEL 7 fails during upgrade of Storage Resource agent](#)
- [You cannot install IBM Spectrum Control on AIX and Linux systems](#)
You cannot install IBM Spectrum Control on AIX and Linux systems. You receive an error message that indicates that the DB2® user ID does not exist.
- [After reboot of the system, IBM Spectrum Control cannot communicate with Db2](#)
After a reboot of the system on Windows, Db2 starts up but IBM Spectrum Control cannot communicate with Db2.
- [Cannot deploy Storage Resource agent because of a timeout, get error message NAD0006E](#)
You are trying to deploy a Storage Resource agent on a system and get this error message.
- [Storage Resource agent information is inconsistent after stop and restart on Windows](#)
On Windows, after you stop the Storage Resource agent, reboot the system, and restart the agent, the information displayed for the agent is not consistent.
- [Cannot deploy Storage Resource agent after changing SSL certificates](#)
You cannot deploy the Storage Resource agent after changing the SSL certificates.
- [Message in Device server log states that port 162 is in use](#)
IBM Spectrum Control uses the default port 162 to listen for SNMP traps. If another process is listening on port 162, you get an error message in the Device server log that port 162 is in use by another process.
- [Unable to log on to the GUI when the user name is defined in both the local OS repository and the Windows domain repository](#)
You are unable to log on to the IBM Spectrum Control GUI with a user name such as "Administrator".
- [BPCIN0057E and BPCIN0072E error messages: Installation on Linux operating system cannot proceed because of issues with DB2 db2inst1 user](#)
In rare occurrences, the IBM Spectrum Control installation program cannot validate the DB2 db2inst1 user during preinstallation on Linux operating systems.
- [Java virtual machine error during uninstall of IBM Spectrum Control](#)
You receive a Java™ virtual machine error message when you cancel an upgrade of IBM Spectrum Control and you try to uninstall the new version of IBM Spectrum Control.
- [CAM-CRP-1613: Regenerating cryptographic keys in Cognos Analytics 11](#)
You encounter a **CAM-CRP-1613** error in the Check password for JVM truststore step when you save a configuration in the IBM® Cognos® Configuration tool. You must regenerate the cryptographic keys in Cognos Analytics 11.
- [BTACS0043E: Probe fails after upgrading IBM Spectrum Control](#)
After upgrading to IBM Spectrum Control 5.3.2 or later, you encounter a **BTACS0043E** error during a probe.

Creating a keystore for an IBM Spectrum Control server

You need this recovery procedure if you do not know the keystore password for the Alert server, Device server, or Web server during an upgrade of IBM Spectrum Control.

Problem

During an upgrade to IBM Spectrum Control 5.3.4 or later, the Alert server, Device server, and Web server keystores are converted from the JKS format to the PKCS12 format. If a particular server's keystore password is not the default value and you do not know the server's current keystore password, then the server's keystore conversion does not occur. You must create a new keystore for that IBM Spectrum Control server before you can proceed with your upgrade of IBM Spectrum Control.

Action

Note: The following procedure assumes that you are running IBM Spectrum Control on a Windows operating system.
To create a keystore, use the following steps:

1. Cancel the IBM Spectrum Control installation program.
2. Stop each IBM Spectrum Control server whose current keystore password you do not know.
3. Make a backup of the following files for each IBM Spectrum Control server whose current keystore password is unknown:
 - `installation_dir\wlp\usr\servers\server_name\resources\security\key.jks`
 - `installation_dir\wlp\usr\servers\server_name\server.xml`where `installation_dir` is the IBM Spectrum Control installation location and `server_name` is the name of the IBM Spectrum Control server. For example, `alertServer`, `deviceServer`, or `webServer`.
4. Delete the following file for each IBM Spectrum Control server whose current keystore password is unknown:

- `installation_dir\wlp\usr\servers\server_name\resources\security\key.jks`
where `installation_dir` is the IBM Spectrum Control installation location and `server_name` is the name of the IBM Spectrum Control server. For example, `alertServer`, `deviceServer`, or `webServer`.
- Open a command window and run the following commands:
 - `cd installation_dir\wlp\bin`
 - `securityUtility.bat createSSLCertificate --server=server_name --password=new_password --validity=825`
 - `securityUtility.bat encode`
When prompted, enter the `new_password` value that you set in Step 5b and make note of the resulting encoded value.

Where `installation_dir` is the IBM Spectrum Control installation location, `new_password` is the new keystore password that you are setting, and `server_name` is the name of the IBM Spectrum Control server. For example, `alertServer`, `deviceServer`, or `webServer`. It is recommended to use the value `default` for the new keystore password.
 - `cd installation_dir\jre\bin`
 - `ikeycmd -cert -delete -db installation_dir\wlp\usr\servers\server_name\resources\security\key.jks -pw new_password -label default`
 - `ikeycmd -cert -create -db installation_dir\wlp\usr\servers\server_name\resources\security\key.jks -pw new_password -label default -size 2048 -sig_alg SHA256_WITH_RSA -expire 825 -dn "CN=<machine_FQDN>, OU=server_name, O=ibm, C=us" -san_dnsname <machine_FQDN> -eku serverAuth`
Where `<machine_FQDN>` is the fully qualified domain name of the machine where you installed the IBM Spectrum Control servers. For example, `myserver.mycompany.com`
 - For each IBM Spectrum Control server whose current keystore password is unknown, edit the `installation_dir\wlp\usr\servers\server_name\server.xml` file, where `installation_dir` is the IBM Spectrum Control installation location and `server_name` is the name of the IBM Spectrum Control server. For example, `alertServer`, `deviceServer`, or `webServer`.
In the `server.xml` file, locate the line for the `keyStore` element (for example - `<keyStore id="home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_troubleshooting_ref_create_new_kstore_servers_defaultKeyStore" password="{xor}OTAwbmw9Pi0=" />`).
Replace the password value with the encoded `new_password` value that you generated in Step 5c. Ensure that the password value is enclosed in double quotation marks and contains the `{xor}` prefix (for example - `<keyStore id="home_markdown_jenkins_workspace_Transform_in_SS5R93_5.4.9_docs_troubleshooting_ref_create_new_kstore_servers_defaultKeyStore" password="{xor}encodedNewPassword" />`).
 - If you upgrade from IBM Spectrum Control 5.3.0 or later and if you create a new keystore for the IBM Spectrum Control Device server, you must use the **keytool** command to update the IBM Spectrum Control data collector trusted certificates or else the data collector does not communicate properly with the IBM Spectrum Control Device server.
Use the following steps, if you **did** create a new keystore for the IBM Spectrum Control Device server.

If you do **not** need to use these steps, proceed to Step 8.
 - `cd installation_dir\jre\bin` where `installation_dir` is the IBM Spectrum Control installation location.
 - Enter the following command to export the SSL certificate from the IBM Spectrum Control new Device server keystore:

`keytool.exe -exportcert -alias default -keystore "installation_dir\wlp\usr\servers\deviceServer\resources\security\key.jks" -storepass new_device_server_keystore_password -file deviceServer.cert`

where `installation_dir` is the IBM Spectrum Control installation location, `new_device_server_keystore_password` is the `new_password` value that you set in Step 5b.
 - Enter the following command to delete the previous IBM Spectrum Control Device server SSL certificate from the IBM Spectrum Control data collector trusted certificates:

`keytool.exe -delete -alias deviceServer -keystore "installation_dir\jre\lib\security\cacerts" -storepass data_collector_keystore_password`

where `installation_dir` is the IBM Spectrum Control installation location and `data_collector_keystore_password` is the IBM Spectrum Control data collector keystore password. The default value for this password is `changeit`.
 - Enter the following command to add the SSL certificate from the IBM Spectrum Control new Device server keystore to the IBM Spectrum Control data collector trusted certificates:

`keytool.exe -importcert -noprompt -trustcacerts -alias deviceServer -file deviceServer.cert -keystore "installation_dir\jre\lib\security\cacerts" -storepass data_collector_keystore_password`

where `installation_dir` is the IBM Spectrum Control installation location and `data_collector_keystore_password` is the IBM Spectrum Control data collector keystore password. The default value for this password is `changeit`.
 - Start each IBM Spectrum Control server that you stopped in Step 2.
 - Start the IBM Spectrum Control installation program again and upgrade to the latest version of IBM Spectrum Control.

Important: If your current IBM Spectrum Control system uses LDAP authentication **with** secure communication between IBM Spectrum Control and the LDAP repository, and if you create a new keystore for the IBM Spectrum Control Web server by using this procedure, you lose the LDAP server's SSL certificate and LDAP authentication does not work. To resolve this situation after you upgrade IBM Spectrum Control to 5.3.4 or later, you must log in to the IBM Spectrum Control GUI with the `tpcFileRegistryUser` user ID and password and complete the *Changing from operating system to LDAP authentication* procedure to reconfigure LDAP authentication with a secure communication between IBM Spectrum Control and the LDAP repository.

If your current IBM Spectrum Control system uses LDAP authentication **without** secure communication between IBM Spectrum Control and the LDAP repository, and if you create a new keystore for the IBM Spectrum Control Web server by using this procedure, LDAP authentication will continue to work after you upgrade to IBM Spectrum Control 5.3.4 or later.

Related tasks

- [Starting and stopping the IBM Spectrum Control servers](#)

- [Changing from operating system to LDAP authentication](#)

Error message: INS3105E Failure to install remote agent

This error occurs when you are installing a remote agent on a UNIX or Linux® system. .

Problem

You get this error message:

```
PM INS3105E: Host <fully_qualified_hostname> is associated
with an invalid loopback IP 127.0.0.1
```

In the `/etc/hosts` file, the line for the specified IP address should be:

```
127.0.0.1    localhost.localdomain localhost
```

and **not**

```
127.0.0.1    <fully_qualified_hostname> <short_hostname>
localhost.localdomain localhost
```

You should have a separate entry for the regular IP address and hostname of the IBM Spectrum® Control server. For example:

```
9.47.98.63 <fully_qualified_hostname> <short_hostname>
```

Action

The DB2® instance might not have stopped properly. Manually stop and restart DB2.

Installing a remote agent fails

This error occurs when you install a remote agent and it fails.

Problem

You are getting a connection error to the remote computer when you try to install a remote agent.

Action

When you are installing a remote agent from a Windows computer, you must have the HOSTS file set properly on the Windows computer. For information about the HOSTS file, see [Changing the HOSTS file](#).

The Solaris HBA identifies a SCSI device instead of fibre channel LUNs

This problem occurs when the Oracle Solaris HBA identifies a SCSI device instead of fibre channel LUNs.

Problem

This problem may be a result of a configuration setup problem with the HBA.

Action

In Solaris, make sure that the `/etc/hba.conf` file points to the correct vendor HBA API library.

For example, if the QLOGIC HBA API library is set as follows in the `/etc/hba.conf` file:

```
ql2x00    libqlsdm.so
```

This is incorrect and should be changed to:

```
ql2x00    /usr/lib/libqlsdm.so
```

In windows, this library is identified using the registry value at:

```
HKEY_LOCAL_MACHINE/SOFTWARE/SNIA/HBA/<vendor_name>
```

For example, for QLogic you will have this value:

```
HKEY_LOCAL_MACHINE/SOFTWARE/SNIA/HBA/QL2X00
```

The property for the library is in the `LibraryFile` key.

For QLogic, a correct configuration has the **LibraryFile** key set as:

```
C:\Program Files\QLogic Corporation\SANsurfer\ql2xhai2.dll
```

Also, the RDAC controller does not have a common HBA API interface. Therefore, the WWN and fibre identification does not work at this time.

Installing IBM Spectrum Control on Windows with installation files on a Samba fileshare

UNIX and Linux® file systems use Samba software to mount file systems on Windows. UNIX and Linux mark hidden files with a leading dot '.' in the file or the folder name.

Problem

In the Samba default configuration, the files and folders with the leading dot get the Windows "hidden" files attribute to ensure consistent behavior.

Action

Change the Samba configuration setting to:

```
hide dot files = no
```

for the Samba share or as a general setting in the Samba configuration file **smb.conf**. The files and folders are no longer marked as hidden.

Note: If your network share is a Samba share on a Linux or AIX® machine, add the following line to the configuration file. This file is in directory /etc/samba/smb.conf.

```
<share_name_where_installation_dir_files_are_located>
.
.
.
#Set to show "dotted" files
hide dot files = no
```

The dots represent any current lines that might be in the share name section.

Restart your Samba service or reload the configuration file. See your operating system documentation for information about restarting the service.

IBM Spectrum Control agent installation fails when files are mapped to a drive in a Windows Terminal Services environment

The IBM Spectrum® Control agent installation fails when files are mapped to a drive in a Windows Terminal Services environment.

Problem

The IBM Spectrum Control agent installation fails when files are mapped to a drive in a Windows Terminal Services environment.

Action

There are two ways to correct this problem:

1. Map the installation source files through the UNC path (for example, \\<server_name>\<disk>) rather than using the drive letter (do not use **net use x:\\<server_name>\<disk>**).
2. Copy the installation files manually to the target host and run the installation program locally.

Db2 services do not start when you restart the system on a UNIX platform

The Db2® services do not start when you perform a system reboot on UNIX.

Problem

On UNIX operating systems, the instance is not enabled for autostart by default.

Action

On UNIX operating systems, to enable an instance to start automatically after each system restart, enter the following command:

```
db2iauto -on <instance_name>
```

where <instance_name> is the login name of the instance.

Cannot install the Storage Resource agent on AIX through RSH

You cannot install the Storage Resource agent on an AIX® system that has the Remote Shell (RSH), but does not have the Secure Shell (SSH).

Problem

The AIX agent system does not have SSH but does have RSH. You can connect with RSH from the server through a command line, but the Storage Resource agent deployment fails with the following message:

```
RSH: ... Could not establish a connection to the target machine using the
credentials supplied
```

Action

Check the `.rhosts` file on the AIX agent system to see if it allows the server and user to connect to AIX through RSH. The `$HOME/.rhosts` file defines which remote hosts (computers on a network) can invoke certain commands on the local host without supplying a password.

The format of the `$HOME/.rhosts` file is:

```
<host_name_field>    [<user_name_field>]
```

Ensure that you add the server and user name to this file and save the file. Deploy the Storage Resource agent again.

Storage Resource agent failed to install using install command

The Storage Resource agent failed to install when using the install command.

Problem

When you specify the install command for the Storage Resource agent, check to see if you specified an agent installation location that contains an ending slash (\). For example, `C:\agent1\`. This causes the installation to fail.

Action

Do not specify an ending slash when you specify the agent installation location. For example, specify `C:\agent1`.

Upgrading to IBM Spectrum Control 5.4.1 on RHEL 7 fails during upgrade of Storage Resource agent

Problem

The upgrade process fails for IBM Spectrum® Control 5.4.1 on RHEL 7 during an upgrade of the Storage Resource agent.

Action

If your upgrade to IBM Spectrum Control 5.4.1 on RHEL 7 fails during an upgrade of the Storage Resource agent, complete the following steps:

1. If you are performing a GUI mode upgrade, click OK in the failure dialog and then cancel out of your upgrade to IBM Spectrum Control 5.4.1.
2. Search the RHEL 7 process table for the term: Agent. Note, there might be one Agent process running from your IBM Spectrum Control installation location. For example - `/opt/IBM/TPC/agent/bin/Agent`. Perform a **kill -9** command on that Agent process.
3. Apply all available updates and patches to your RHEL 7 operating system.
4. Resume your upgrade to IBM Spectrum Control 5.4.1.

You cannot install IBM Spectrum Control on AIX and Linux systems

You cannot install IBM Spectrum® Control on AIX® and Linux® systems. You receive an error message that indicates that the DB2® user ID does not exist.

Problem

You receive this error message:

```
User ID db2inst1 does not exist on the system
```

You receive this error message even though the db2inst1 user ID exists and DB2 is working properly. The error occurs because the shadow util package is not installed and configured on the system.

A shadow password file is a system file in which encrypted user passwords are stored. Normally, user information for each user is stored in an encrypted format in `/etc/passwd`. However, for shadow passwords, the shadow password file is moved to a separate database for local files, usually in `/etc/shadow` on AIX and Linux

systems.
The shadow password file can only be read by the root user.

Action

Check to see that the shadow passwords are installed and configured on the system.

If the file `/etc/shadow` exists on the system, then the shadow passwords are already installed and configured.

If not, install and configure the shadow passwords. Install the shadow utils package from the operating system CD.

For example, on Linux follow these steps:

1. Install the shadow utils package, shadow-utils-<version_release_architecture>.rpm.
2. To configure the shadow passwords after you install the package, run the following commands:

```
/usr/sbin/pwconv  
/usr/sbin/grpconv
```

Note: The instructions and package name and version might differ between operating systems. Also, the instructions for how to configure the shadow passwords might differ. Check the operating system CD or your vendor for instructions on how to install and configure the shadow passwords.

After reboot of the system, IBM Spectrum Control cannot communicate with Db2

After a reboot of the system on Windows, Db2® starts up but IBM Spectrum® Control cannot communicate with Db2.

Problem

To determine if this problem exists, open a command prompt window and run the command **netstat -na**. Port 25000 is not listed anywhere as a result of this command.

Action

If port 25000 is displayed, follow these steps to enable the Db2 service port for remote connections:

1. Open the Windows service file:

```
C:\WINDOWS\system32\drivers\etc\services
```

2. Search for the following line:

```
db2c_DB2 25000/tcp
```

Note the port number (the port number is probably 25000).

3. Open a Db2 command prompt window by running **db2cmd**. Run the following commands:

```
db2 update dbm cfg using SVCENAME db2c_DB2  
db2set  
db2stop  
db2start
```

Note:

- When you run the **db2set** command, an example of the output is:

```
DB2COMM=tcPIP
```

- If `db2c_DB2` is not the service name of the Db2 instance name on the system, use the name listed in the service file.
- If on a subsequent reboot of the system, this problem occurs again, then Db2 is not storing this information in the configuration file. Try running these commands:

```
db2 update dbm cfg using SVCENAME db2c_DB2 deferred  
(This forces DB2 to update the configuration file.)  
db2set  
db2stop  
(If DB2 displays an error message saying that the database manager was  
not stopped because the databases are still active, use the command  
db2stop force.)  
db2start
```

When you run the **db2set** command, an example of the output is:

```
DB2COMM=tcPIP
```

- Determine if the problem is fixed by running **netstat -na**.
- Verify that port 25000 is listed. For example, your output looks like this:

```
TCP 0.0.0.0:25000 0.0.0.0:0 LISTENING
```

As long as that port is active, Db2 can connect to IBM Spectrum Control.

Cannot deploy Storage Resource agent because of a timeout, get error message NAD0006E

You are trying to deploy a Storage Resource agent on a system and get this error message.

Problem

When you deploy a Storage Resource agent on a system, you see this error message in the log file:

```
NAD0006E Exception thrown for method updateNAOsInfo:
java.net.ConnectException: CTGRI0081E RXA internal
command run on xxxx.xxxx.xxxx.xxxx could not be completed
during the specified timeout interval.
```

You cannot deploy the Storage Resource agent.

Action

Use the tpctool command line interface to increase the default_timeout value for the Storage Resource agent.

1. Use the getdscfg command to see the current values for the Storage Resource agent:

```
tpctool> getdscfg -context SRA
```

The following output is returned:

Property	Context	Value
default_timeout	SRA	120000
long_timeout	SRA	300000
nfs_timeout	SRA	10000

2. Use the setdscfg command to change the default_timeout value:

```
tpctool> setdscfg -context SRA -property default_timeout 150000
```

3. Use the getdscfg command again to verify the new value for the default_timeout property:

```
tpctool> getdscfg -context SRA
```

The following output is returned:

Property	Context	Value
default_timeout	SRA	150000
long_timeout	SRA	300000
nfs_timeout	SRA	10000

Storage Resource agent information is inconsistent after stop and restart on Windows

On Windows, after you stop the Storage Resource agent, reboot the system, and restart the agent, the information displayed for the agent is not consistent.

Problem

For example, the IP address and agent name are not the same before and after you stop and start the Storage Resource agent.

Action

If the IBM Spectrum® Control server and Storage Resource agent (or Fabric agent) are installed on the same Windows system that is also the Primary Domain Controller and has Microsoft Failover Cluster node, then the hosts file must be changed.

You must enter the correct IP address, fully qualified domain name, and host name in the hosts file before installing IBM Spectrum Control. The hosts file is found in this directory:

```
C:\Windows\system32\drivers\etc\hosts
```

Enter a line in the hosts file, for example:

```
98.256.43.23 xyz.company.com xyz
```

If the IP address, fully qualified domain name, and host name are not added to the hosts file, then the preferred private IP address (169.x.y.z) for the Microsoft Failover Cluster Virtual Adapter is incorrectly returned as the IP address for the Data server and Storage Resource agent.

Cannot deploy Storage Resource agent after changing SSL certificates

You cannot deploy the Storage Resource agent after changing the SSL certificates.

Problem

When generating custom SSL certificates, the certificates have a start date, end date, and time when they are valid. These dates and times are related to the system where these custom certificates were generated (which is usually the server system). When installing a Storage Resource agent on a remote system, you must check the date and

time on the Storage Resource agent system. If the server and agent systems are in the same time zone, they must have the same date and time. Otherwise, the time zone difference should be set.
For example, if the server system is 8:00 PM, the agent system should also be 8:00 PM. If the agent system is set at a different time (for example, 6:00 PM) at the time the SSL custom certificates are generated on the server system with a time of 8:00 PM, the deployment of the Storage Resource agent will fail.

Action

Set the time zones on the agent system to match the server time.

Message in Device server log states that port 162 is in use

IBM Spectrum® Control uses the default port 162 to listen for SNMP traps. If another process is listening on port 162, you get an error message in the Device server log that port 162 is in use by another process.

Problem

You see this message in the Device server message log if another process is listening on port 162. IBM Spectrum Control uses the default port 162 to listen for SNMP traps. The Windows SNMP Trap service also uses port 162 to listen for SNMP traps.

Action

If the Windows SNMP Trap service is running, you can stop the service or change the IBM Spectrum Control default port. To change the default port number for IBM Spectrum Control, use the **tpctool** command line interface (CLI).

To set the port number for the Device server, complete the following steps:

1. Run the **getdsCfg** command to determine the current value of the port. From the command prompt, enter the following command:

```
tpctool getdsCfg -user username -pwd password -url host:port
-property SnmpTrapPort
```

Where:

- *user* is a IBM Spectrum Control user ID.
- *password* is the password for the IBM Spectrum Control user ID.
- *host* is the host name or IP address and *port* is a valid port number for the HTTP service of the Device server.

2. Run the **setdsCfg** command to change the port. Run the following command:

```
tpctool setdsCfg -user username -pwd password -url host:port
-property SnmpTrapPort -context user port_value
```

Where:

- *user* is a IBM Spectrum Control user ID.
- *password* is the password for the IBM Spectrum Control user ID.
- *host* is the host name or IP address and *port* is a valid port number for the HTTP service of the Device server.
- *port_value* is the port that you specify for IBM Spectrum Control to listen for SMNP traps.

Note: **-context user** must be entered exactly as shown in the preceding example.

3. Run the **getdsCfg** command again to verify the change to the port number.

For more information about **tpctool**, see [tpctool command](#). You can also view help from the command line by entering the command with the **-help** option.

Unable to log on to the GUI when the user name is defined in both the local OS repository and the Windows domain repository

You are unable to log on to the IBM Spectrum® Control GUI with a user name such as "Administrator".

Problem

You are unable to log on to the IBM Spectrum Control GUI with a user name such as "Administrator" when the user name is defined in both the local operating system repository and the Windows domain repository.

Cause

This problem occurs because a user name can have some IBM Spectrum Control roles mapped to local operating system groups and other roles to Windows domain groups.

Environment

This problem occurs on Windows operating systems.

Resolving the problem

If a user name, such as "Administrator", is defined in both the local operating system repository and the Windows domain repository, you must use a fully qualified name that contains either a domain name prefix or a machine prefix to log on to the IBM Spectrum Control GUI, for example, *domain_name*\Administrator or *machine_name*\Administrator.

Related tasks

- [Authorizing users](#)

BPCIN0057E and BPCIN0072E error messages: Installation on Linux operating system cannot proceed because of issues with DB2 db2inst1 user

In rare occurrences, the IBM Spectrum® Control installation program cannot validate the DB2® db2inst1 user during preinstallation on Linux® operating systems.

Problem

In this situation, the following messages are generated during a preinstallation check and are displayed at the beginning of the installation process:

BPCIN0057E The user name db2inst1 is not in an operating system group that has DB2 SYSADM authority.

BPCIN0072E The validation for user name db2inst1 has failed. Check to see if this user name exists or if DB2 is running.

Action

Determine whether the DB2 db2inst1 user has SYSADM authority and whether DB2 is running.

If the db2inst1 user does not have SYSADM authority or DB2 is not running, correct the problem and then continue the installation.

If the user has the correct authority and DB2 is running, cancel the installation and start the installation program again.

Java virtual machine error during uninstall of IBM Spectrum Control

You receive a Java™ virtual machine error message when you cancel an upgrade of IBM Spectrum® Control and you try to uninstall the new version of IBM Spectrum Control.

Problem

You upgrade IBM Spectrum Control to a new version. However, IBM Spectrum Control does not start after the upgrade. You attempt to uninstall the upgrade and you get a No Java virtual machine could be found error.

Solution

The Java virtual machine directory was renamed during the upgrade from `USER_INSTALL_DIR/jre` to `USER_INSTALL_DIR/jre.bck`. Rename it back to `USER_INSTALL_DIR/jre`.

CAM-CRP-1613: Regenerating cryptographic keys in Cognos Analytics 11

You encounter a **CAM-CRP-1613** error in the Check password for JVM truststore step when you save a configuration in the IBM® Cognos® Configuration tool. You must regenerate the cryptographic keys in Cognos Analytics 11.

Problem

When you complete the upgrade of your Cognos Analytics 11 environment, you save the configuration. You might encounter a failure in the Check password for JVM truststore step in the IBM Cognos Configuration tool. The following error is displayed in the Details area in IBM Cognos Configuration:

CAM-CRP-1613 Not able to retrieve configured password

Action

You must regenerate cryptographic keys in Cognos Analytics 11. For more information about how to regenerate cryptographic keys, see [How to regenerate cryptographic keys in Cognos Analytics 11](#).

BTACS0043E: Probe fails after upgrading IBM Spectrum Control

After upgrading to IBM Spectrum® Control 5.3.2 or later, you encounter a **BTACS0043E** error during a probe.

Problem

After you complete the upgrade of IBM Spectrum Control to 5.3.2 or later, and a probe attempts to collect metadata from a device, authentication to the IBM Spectrum Control server fails.

Error **BTACS0043E** is shown in the following log files:

Data collection log

```
STS0306I: Job queued for processing. Waiting for idle thread.  
GEN0404I: Probe started.  
SAA0022I: Storage Subsystem XYZ (0) will be probed.  
SRV0451E: Error calling device server probeAll service.  
BTACS0043E Failed to authenticate with host my.spectrum.control.server.ibm.com:9550.  
Invalid host authentication password.  
GEN0400E: Probe completed with errors.
```

Server log

```
GEN0042E: Error reading log-file installation_dir/device/log\SVC-2145-0123456789-IBM.340\msg.33  
520527.1.0000020012345678.log  
BTACS0043E Failed to authenticate with host my.spectrum.control.server.ibm.com:9550.  
Invalid host authentication password.  
java.lang.RuntimeException: BTACS0043E Failed to authenticate with host my.spectrum.control.server.ibm.com:9550:. Invalid host  
authentication password.
```

Action

Contact IBM® Software Support for assistance. For more information about this issue, see <https://www.ibm.com/support/pages/apar/IT31727>.

Related reference

- [Default locations of log files](#)

SAN Volume Controller troubleshooting

Use this section to troubleshoot and resolve SAN Volume Controller problems.

- [All ports of a host lose access to the volume](#)
This problem occurs with SAN Volume Controller.
- [Error message: performance data files could not be correlated](#)
This error occurs with the SAN Volume Controller.
- [SAN Volume Controller reports are incorrect](#)
The SAN Volume Controller reports are incorrect.
- [Probing a SAN Volume Controller cluster fails when the cluster ID changes](#)
When a SAN Volume Controller cluster ID changes, the first probe after the ID change fails. Subsequent probes complete successfully.
- [Setting the rate at which data is copied between volumes on storage systems that run IBM Spectrum Virtualize](#)
Set the synchronization rate to avoid performance problems on volumes on IBM® SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, and IBM Storwize® storage systems, and on IBM FlashSystem® devices that run IBM Spectrum Virtualize.
- [Some capacity values are zero](#)
Storage systems that run IBM Spectrum® Virtualize with firmware 8.2 or earlier might show zero values for some capacity values.

All ports of a host lose access to the volume

This problem occurs with SAN Volume Controller.

Problem

Unassigning one or more ports of a host unassigns all ports of the host.

If you use the SAN Volume Controller user interface or the SAN Volume Controller CLI to create a host that has several ports, and then unassign one or more ports of a host using IBM Spectrum® Control, all ports of the host will lose access to the volume.

Action

Do not create hosts with several ports.

Error message: performance data files could not be correlated

This error occurs with the SAN Volume Controller.

Problem

For SAN Volume Controller subsystems, this error message can occur during performance data collection.

Action

This is usually a problem with mismatching timestamps for the I/O statistics dump files on the SAN Volume Controller itself. This problem can be resolved by setting the time zone properly on the cluster.

SAN Volume Controller reports are incorrect

The SAN Volume Controller reports are incorrect.

Problem

An example of an incorrect report is when the report shows 0 for the allocated volume space when the report should show a value.

Action

When you have hosts connected to storage subsystems that have multipathing enabled, install the multipathing subsystem device drivers (SDD) on the hosts.

Probing a SAN Volume Controller cluster fails when the cluster ID changes

When a SAN Volume Controller cluster ID changes, the first probe after the ID change fails. Subsequent probes complete successfully.

Tip: When a SAN Volume Controller cluster is rebuilt or recovered, a new cluster ID is generated. To ensure that the first probe after a cluster ID change is successful, complete the following steps.

1. Remove the cluster from IBM Spectrum® Control before you rebuild or recover the cluster.
2. Rebuild or recover the cluster.
3. Add the cluster to IBM Spectrum Control.

Setting the rate at which data is copied between volumes on storage systems that run IBM Spectrum Virtualize

Set the synchronization rate to avoid performance problems on volumes on IBM® SAN Volume Controller, IBM Spectrum Virtualize for Public Cloud, IBM Spectrum Virtualize as Software Only, and IBM Storwize® storage systems, and on IBM FlashSystem® devices that run IBM Spectrum Virtualize.

The synchronization rate on storage systems that run IBM Spectrum Virtualize is the rate at which the volume copies synchronize again after synchronization is lost. IBM Spectrum® Control can do tasks that significantly alter the structure of storage. For example, IBM Spectrum Control can add a mirrored volume copy, or migrate volumes from one pool to another by using the storage optimizer functions. When IBM Spectrum Control does these tasks, by default, the synchronization rate is set to 50%.

The synchronization rate can cause performance problems on all volumes on storage systems that run IBM Spectrum Virtualize. However, you can use the **tpctool** command to change the synchronization rate.

To verify the current synchronization rate, go the installation_dir\cli directory, then use the **getdscfg** command:

```
tpctool getdscfg -user admin_user_id -pwd admin_user_password -url localhost:9550
-property OptAutoController.SyncRate -context DeviceServer
```

If the synchronization rate is set to the default value of 50%, the **tpctool** command shows the following information:

Property	Context	Value
=====		
OptAutoController.SyncRate	DeviceServer	50

To set the value of the synchronization rate to 70%, use the **setdscfg** command:

```
tpctool setdscfg -user admin_user_id -pwd admin_user_password -url localhost:9550
-property OptAutoController.SyncRate -context DeviceServer 70
```

The new synchronization rate is used for all storage systems that run IBM Spectrum Virtualize that you add to IBM Spectrum Control. You cannot set this value independently for each resource.

To verify that the value of the synchronization rate is modified, use the **getdscfg** command again. If the synchronization rate is set to 70%, the **tpctool** command shows the following information:

Property	Context	Value
=====		
OptAutoController.SyncRate	DeviceServer	70

After you modify the value of the synchronization rate, stop and restart the Device server.

Related reference

- [setdscfg](#)

Some capacity values are zero

Storage systems that run IBM Spectrum® Virtualize with firmware 8.2 or earlier might show zero values for some capacity values.

Problem

This issue occurs because the capacity values for these storage systems were changed to rely on physical capacity values.

Action

Upgrade the firmware for IBM® Storage Virtualize to a newer version. This issue is resolved in newer versions of firmware, starting with the following versions:

- 8.1.3.6
- 8.2.1.4

For more information, see https://public.dhe.ibm.com/storage/san/sanvc/release_notes/813_releasenotes.html#APARs.

DS8000 troubleshooting

Use this section to troubleshoot and resolve DS8000® problems.

- [Host connections for DS8000 storage systems are not being displayed](#)
Host connections that were created by using the **mkhostconnect** command might not be displayed on the details page for DS8000 storage systems.
- [The user account is locked after the DS Storage Manager password for the HMC is changed](#)
If the DS Storage Manager password for the Hardware Management Console (HMC) is changed, you must update the password in IBM Spectrum® Control.
- [Get warning messages STA0044I, STA0035W, and STA0036I](#)
These warning messages occur when you perform an agent probe.

Host connections for DS8000 storage systems are not being displayed

Host connections that were created by using the **mkhostconnect** command might not be displayed on the details page for DS8000® storage systems.

Problem

On a DS8000 storage system, host connections that were created by using the **mkhostconnect** command might not be displayed on the details page for that storage system in IBM Spectrum® Control. Hosts that were configured with the **mkhost** command are displayed properly.

Action

To ensure that the host connections that were configured with **mkhostconnect** (and any related agentless servers) are displayed properly, complete the following steps:

1. Go to <https://www.ibm.com/support/pages/node/690153>.
2. To migrate the host connections, follow the process in the section called "The GUI Migration Process" on page 10.
3. In IBM Spectrum Control, schedule a probe to collect the latest metadata from the DS8000 storage system. When the probe is complete, the host connections for the storage system are displayed properly on its details page.

Tips:

- In the future, when you configure host connections in DS8000, use the **mkhost** command to ensure that the connections are shown in IBM Spectrum Control.
- In cases where multiple agentless servers are created for host connections that are on one server, you can multi-select those servers on the Servers page in IBM Spectrum Control and then right-click to merge them into a single agentless server.
- z/OS® hosts that use storage volumes from a DS8000 are not displayed as host connections in IBM Spectrum Control. This behavior occurs because z/OS hosts are not configured by using **mkhostconnect** or **mkhost**.

For more information about the changes for host configuration in the DS8000 Release 7.4 Storage Management GUI, see <https://www.ibm.com/support/docview.wss?uid=s8g1S1005012>.

The user account is locked after the DS Storage Manager password for the HMC is changed

If the DS Storage Manager password for the Hardware Management Console (HMC) is changed, you must update the password in IBM Spectrum® Control.

Problem

If you do not update this password in IBM Spectrum Control, the DS Storage Manager user account for the HMC might be locked when you attempt to connect to the storage system with the incorrect password.

Action

Update the password in the IBM Spectrum Control GUI as shown in the following table.

Interface	Action
IBM Spectrum Control	<p>Update the password for the storage system as shown in the following steps:</p> <ol style="list-style-type: none">1. In the menu bar, go to Storage > Block Storage Systems.2. Right-click the storage system that you want to update and click Connections > Modify Connection.3. In the Enter User Credentials window, click User Name and Password in the Authentication field.4. Enter the user name and password.5. Click OK.
IBM Spectrum Control	<p>This task is required only if both of the following conditions are true:</p> <ul style="list-style-type: none">• You are using the launch-in-context feature to start the element manager, DS Storage Manager, from IBM Spectrum Control.• You are not using single sign-on with LDAP authentication to connect to the element manager. LDAP authentication is required for IBM Spectrum Control to connect to System Storage® DS8000® 4.2 and later. <p>Update the password for the element manager as shown in the following steps. If you have multiple users who are accessing the element manager and their password was changed, you must update the password for each user.</p> <ol style="list-style-type: none">1. In the menu bar, go to Storage > Block Storage Systems.2. Right-click the storage system that you want to update and select Open Storage System GUI.3. Log on to IBM System Storage DS8000 Storage Manager and update the user password. Follow the instructions that are provided in product documentation at https://www.ibm.com/docs/en/ds8900/9.3.0.

Get warning messages STA0044I, STA0035W, and STA0036I

These warning messages occur when you perform an agent probe.

Problem

If a job log of an agent probe shows warnings similar to the following messages:

```
3/15/07 9:22:10 AM STA0044I: Reading capacity
3/15/07 9:22:10 AM STA0035W: Error reported by
device \\.\PhysicalDrive24
3/15/07 9:22:10 AM STA0036I: Sense code: xxxx
                        ASC:      xxxx
                        ASCQ:     xxxx
```

Action

There may be a problem with the operating system or the environment. The corresponding volumes should be checked.

Independent software vendors troubleshooting

Use this section to troubleshoot and resolve independent software vendor problems.

- [HDS HiCommand shows a different set of volumes](#)
This problem occurs with the Hitachi Data Systems.
- [Reports for an Hitachi Data Systems subsystem do not show current information](#)
The IBM Spectrum® Control reports do not show the current information for an Hitachi Data Systems (HDS) subsystem after the HDS Storage Navigator is used to create and assign LUNs.
- [Tool to dump MSCS cluster configuration](#)
This topic provides information on the dump tool for the MSCS cluster configuration.
- [Cannot collect performance data from EMC CLARiiON](#)
You cannot collect performance data from the EMC CLARiiON.
- [Error message: HWNPM2132W](#)
You receive this message when you are try to collect performance data from a Data ONTAP file system.
- [Assigning a volume for EMC CLARiiON fails](#)
If you try to assign a volume to a host for an EMC CLARiiON storage system and that volume is not defined on the EMC CLARiiON storage system, the assignment fails.
- [Incorrect password for SAN Volume Controller is accepted](#)
A valid password with extra characters is accepted when you add a SAN Volume Controller storage system to IBM Spectrum Control .
- [You cannot unassign volumes from a Hewlett Packard EVA device](#)
IBM Spectrum Control cannot unassign volumes from a Hewlett Packard EVA device.
- [Getting timeouts and probe failures when using a Brocade switch](#)
You are getting IBM Spectrum Control timeouts and probe failures when using a Brocade switch.
- [Resolving a problem connecting to ESX](#)
If a problem occurs with the connection to the VMware ESX hypervisor, the messages BTAVM2201E/BTAVM2202E, BTAVM2017E, and BTADS0001I are displayed in the probe result and trace logs. This message sequence indicates that the optical-fiber connector (FC) connection to the hypervisor was disabled .

HDS HiCommand shows a different set of volumes

This problem occurs with the Hitachi Data Systems.

Problem

The HDS **HiCommand** shows a different set of volumes for an Hitachi Data System subsystem from what is reported by IBM Spectrum® Control. This is a known problem with either the **HiCommand** CIMOM or the **HiCommand** GUI.

Action

Contact Hitachi Data Systems customer support to see if a fix is available.

Reports for an Hitachi Data Systems subsystem do not show current information

The IBM Spectrum® Control reports do not show the current information for an Hitachi Data Systems (HDS) subsystem after the HDS Storage Navigator is used to create and assign LUNs.

Problem

This problem occurs after you have used the HDS Storage Navigator to create and assign LUNs, add the HDS SMI-S agent, and then run a discovery and probe job.

Action

To get current information, follow these steps:

1. Go to the HDS HiCommand Device Manager that contains the SMI-S agent. Click **Resources > Subsystems** in the navigation tree.
2. Click on the HDS subsystem you want to refresh. You will see the **Refresh** button under Subsystem IP. Click **Refresh**.
3. Probe the HDS subsystem in IBM Spectrum Control.
4. You can now look at your reports for the HDS subsystem.

Tool to dump MSCS cluster configuration

This topic provides information on the dump tool for the MSCS cluster configuration.

Problem

When looking at an MSCS cluster problem, the current state of the clustering environment is often needed. Microsoft provides a tool to dump the MSCS cluster configuration into a set of files.

Action

For more information about the tool to dump the MSCS cluster configuration, search for the scripted system configuration gathering tools at <http://www.microsoft.com>

Cannot collect performance data from EMC CLARiiON

You cannot collect performance data from the EMC CLARiiON.

Problem

The EMC CIM agent incorrectly reports that it supports BSP even when the CIM agent has not been enabled to run the BSP service.

Action

To enable the CIM agent for performance data, follow these steps:

1. Run the **testsmiprvider** command to discover the Symmetrix devices.
 - a. Open a Command Prompt window. Go to the EMC directory:

```
<EMC_SMIS_Provider_base_dir>\symcli\storbin
```

Example:

```
C:\Program Files\EMC\SYMCLI\storbin
```

- b. Run the **testsmipprovider** command in the Command Prompt window. Provide the information for the Host, Connection Type, Port, Username, and Password fields. Use the default values unless you have changed them.
 - c. When you get to the Command Prompt window, run the **disco** command to discover the Symmetrix devices. The command runs and ends with an output of 0 (success). It does not show any discovered devices at this point.
 - d. Press the **Enter** key and then **q** to exit.
2. Use the **symcfg list** command to list the Symmetrix devices that have been discovered. This command displays each Symmetrix device that the **testsmipprovider** command discovered.
3. Register the EMC SMI-S provider in IBM Spectrum® Control as a CIM agent. Use the following values:

Host

Host name or IP address of the EMC SMI-S provider host.

Port and Protocol

- 5988 (unsecure - http)
- 5989 (secure - https)

Userid/password

The EMC SMI-S provider user ID and password that was defined when you configured the SMI-S provider. If user authentication is not enabled on the CIM agent, enter any values you want.

Description of the CIM Agent

A value that is used to identify the managed device or devices.

4. Run an IBM Spectrum Control CIMOM discovery.
5. To enable performance monitoring, run the **symcfg** command. For example:

```
symcfg authorization add -host <HostName>  
-username <UserName> -password <password>
```

6. Run the command in step 5 for each of the EMC controllers.

Error message: HWNPM2132W

You receive this message when you are try to collect performance data from a Data ONTAP file system.

Problem

The Data ONTAP SMI-S 3.0 agent has the following limitation: If a volume is offline, the performance monitor might fail with the message: **PM HWNPM2132W Performance data could not be collected for device <device>.**

Action

Check your vendor's documentation for any current information in this area.

Assigning a volume for EMC CLARiiON fails

If you try to assign a volume to a host for an EMC CLARiiON storage system and that volume is not defined on the EMC CLARiiON storage system, the assignment fails.

Problem

You cannot assign a volume to a host for an EMC CLARiiON storage system.

Action

You must manually define the host node WWN on the EMC CLARiiON storage system before you can complete the host port assignment on IBM Spectrum® Control.

Incorrect password for SAN Volume Controller is accepted

A valid password with extra characters is accepted when you add a SAN Volume Controller storage system to IBM Spectrum® Control .

Problem

The valid password with the extra characters should not be accepted. The entire password should be checked and not just the first eight characters.

Action

Enter an acceptable password.

You cannot unassign volumes from a Hewlett Packard EVA device

IBM Spectrum® Control cannot unassign volumes from a Hewlett Packard EVA device.

Problem

You cannot unassign volumes from a Hewlett Packard EVA device.

Action

To unassign volumes on this device, use the Hewlett Packard EVA element manager.

Getting timeouts and probe failures when using a Brocade switch

You are getting IBM Spectrum® Control timeouts and probe failures when using a Brocade switch.

Problem

When a Brocade i10000 switch has a faulty control processor card, the SMIS agent managing the switch can hang indefinitely. This issue causes timeouts and probe failures from the IBM Spectrum Control server.

Action

To resolve the issue, contact your switch provider.

Resolving a problem connecting to ESX

If a problem occurs with the connection to the VMware ESX hypervisor, the messages BTAVM2201E/BTAVM2202E, BTAVM2017E, and BTADS0001I are displayed in the probe result and trace logs. This message sequence indicates that the optical-fiber connector (FC) connection to the hypervisor was disabled .

Problem

When you run a probe on a VMware ESX hypervisor, you might receive the following sequence of error messages:

- Message **BTAVM2201E** indicates that the physical storage configuration of the hypervisor could not be collected.
- Message **BTAVM2202E** indicates that the logical storage configuration of the hypervisor could not be collected.
- Message **BTAVM2017E** indicates that the probe of the hypervisor failed.
- Message **BTADS0001I** indicates that the specified discovery request has completed, but with a return code of **99** The probe has failed.

The following example shows the error sequence, as generated on an VMware ESX hypervisor, version 5:

```
2012-04-10 05:10:01.883-0700 BTADS0000I Starting Discover Process Front End ,
with Device Server RUN ID 1025 , and Job ID 2182 .
2012-04-10 05:10:01.890-0700 BTAVM0008I Probe of hypervisor
tpcblade6-4.storage.server.domain.com has started.
2012-04-10 05:10:09.316-0700 BTAVM0015I Collection of the
physical storage configuration for hypervisor tpcblade6-4.storage
.server.domain.com has started.
2012-04-10 05:20:10.521-0700 BTAVM2011E The operation
performMethodInvocation could not complete within the time limit of
600000 milliseconds.
2012-04-10 05:20:10.564-0700 BTAVM2201E Probe: An error occurred
during the collection of the physical storage configuration.
2012-04-10 05:20:10.597-0700 BTAVM2017E Probe of the hypervisor
tpcblade6-4.storage.server.domain.com failed.
2012-04-10 05:20:10.607-0700 BTADS0001I Discover Process with
Device Server RUN ID 1025 and Job ID 2182 is complete with
Status= 4 , Return Code= 99 .
```

The point of failure can vary between different VMware ESX hypervisor versions. For example, a VMware ESX hypervisor, version 3 will fail during the collection of logical storage with **BTAVM2202E**, and a VMware ESX hypervisor, version 5 fails during the collection of physical storage with **BTAVM2201E**.

If messages of this general type are in the probe and the IBM Spectrum® Control device server trace, it is likely that the optical-fiber connector (FC) connection to the hypervisor has been disabled.

The device server trace log is located in the *TPC_Directory/device/log/traceTPCDeviceServer.log*. If this error condition has occurred, the trace log contains an entry similar to the following example:

```
2012-03-07 10:27:34.900-0800 Occured Exception on
https://mystorage.server.com:443/sdk:
com.ibm.tpc.vmmgr.client.VIException: InvocationTargetException.
    at com.ibm.tpc.vmmgr.client.VIClientWrapper.handleException
(VIClientWrapper.java:739)
```

Action

To resolve the problem, reconnect the optical-fiber connector (FC) connection, run the probe again, and refresh the data sources:

1. Add the ESX hypervisor to IBM Spectrum Control to monitor and discover.
2. Probe the ESX hypervisor.
3. Disable the switch port in which the ESX hypervisor is connected.
4. Unplug the cable from the switch or the ESX hypervisor.
5. Run a probe on the ESX hypervisor. The expected result is that the probe fails.
6. Reestablish the connection, either by enabling the port or by reconnecting the FC cable.
7. Scan the HBA on the ESX hypervisor and refresh the datastore objects on the hypervisor using the VMware Infrastructure (VI) client.
8. Run a probe on the ESX hypervisor.

Language troubleshooting

Learn how to troubleshoot and resolve problems with non-English versions of the product.

- [Changing the language in a web browser doesn't change the entire IBM Spectrum Control GUI](#)
When one of the IBM Spectrum® Control internal components is running in a language other than English, you can't change the IBM Spectrum Control GUI to English. However, it can be changed to any language other than English (the default language), although some elements in the GUI might be displayed in an incorrect language.
- [PDF files do not export correctly in languages and locations with non-Latin 1 code pages](#)
In the IBM Spectrum Control, PDF files do not export correctly in languages and locations where non-Latin 1 code pages are used.
- [Traditional Chinese characters not displayed correctly in GUI help panels](#)
This problem occurs when you are displaying Traditional Chinese characters in the IBM Spectrum Control GUI help panels.
- [Cannot read Korean language logs from GUI](#)
This problem occurs in Korean language environments.

Changing the language in a web browser doesn't change the entire IBM Spectrum Control GUI

When one of the IBM Spectrum® Control internal components is running in a language other than English, you can't change the IBM Spectrum Control GUI to English. However, it can be changed to any language other than English (the default language), although some elements in the GUI might be displayed in an incorrect language.

About this task

To display all GUI elements in the correct language, you must update the Java™ virtual machine (JVM) configuration options for the Web server and the Alert server.

Procedure

To modify the internal component to English, complete the following steps:

1. Stop the IBM Spectrum Control Web server and the Alert server.
For instructions on how to stop servers, see [Stopping the IBM Spectrum Control servers by using scripts](#).
2. Update the JVM language and region options for the Web server:
 - a. Open a command line and go to *installation_dir/wlp/usr/servers/webServer*, where *installation_dir* is the directory where IBM Spectrum Control is installed.
 - b. Edit the *jvm.options* file and add these lines:
 - `-Duser.language=en`
 - `-Duser.region=GB`
 - c. Save the file.
3. Update the JVM language and region options for the Alert server:
 - a. Open a command line and go to *installation_dir/wlp/usr/servers/alertServer*, where *installation_dir* is the directory where IBM Spectrum Control is installed.
 - b. Edit the *jvm.options* file and add these lines:
 - `-Duser.language=en`
 - `-Duser.region=GB`
 - c. Save the file.
4. Start the IBM Spectrum Control Web server and the Alert server.
For instructions on how to start servers, see [Starting the IBM Spectrum Control servers by using scripts](#).
5. Open a web browser and set its language to English.
For instructions on how to change the language in a web browser, see [Changing the language of the IBM Spectrum Control GUI](#).
6. Optional: Confirm the changes by viewing the log files.
 - a. Open a command line and go to *installation_dir/wlp/usr/servers/webServer/logs*.
 - b. Open *console.log* and confirm that its content is in English.
 - c. Open a command line and go to *installation_dir/wlp/usr/servers/alertServer/logs*.
 - d. Open *console.log* and confirm that its content is in English.

PDF files do not export correctly in languages and locations with non-Latin 1 code pages

In the IBM Spectrum® Control, PDF files do not export correctly in languages and locations where non-Latin 1 code pages are used.

Problem

When you export information from a page in the IBM Spectrum Control GUI to a PDF, the non-Latin 1 characters are stripped from the output. The following languages are affected:

- Chinese (Simplified)
- Chinese (Traditional)
- Czech
- Hungarian
- Japanese
- Korean
- Polish
- Russian

When you export information to CSV and HTML, the correct output is generated in all languages and locations.

Action

You must use IBM® Cognos® Analytics to generate reports and export these reports as PDF files. For more information about exporting reports as PDF files, see [IBM Cognos Business Intelligence knowledge documentation Version 11.2.0](#)

Traditional Chinese characters not displayed correctly in GUI help panels

This problem occurs when you are displaying Traditional Chinese characters in the IBM Spectrum® Control GUI help panels.

Problem

The Traditional Chinese characters are not displayed correctly in the IBM Spectrum Control GUI help panels.

Action

To work around this problem, follow these steps:

1. Download the file `mtsans_t.zip` from the following Web site: ftp://submit.boulder.ibm.com/download/typography/fonts/worldtype/archive/wts_/2002-02-26/.
2. Unzip the file `mtsans_t.zip` to access the unicode font file `mtsans_t.ttf` (Monotype Sans WT TC).
3. Install the font file `mtsans_t.ttf` on the Traditional Chinese Windows system.
4. Modify the file:

```
<TPC_install_dir>\jre\lib\fontproperties.zh_TW
```

Add the line:

```
filename.Monotype_Sans_WT_TC=MTSANS_T.TTF
```

Modify the `fontproperties.zh_TW` file with the following lines in bold highlight:

```
===== fontproperties.zh_TW =====
serif.0=Times New Roman
serif.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
serif.2=Lucida Sans Regular
serif.3=Times New Roman WT TC
serif.4=EUDC

serif.italic.0=Times New Roman Italic
serif.italic.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
serif.italic.2=Lucida Sans Oblique
serif.italic.3=Times New Roman WT TC
serif.italic.4=EUDC

serif.bold.0=Times New Roman Bold
serif.bold.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
serif.bold.2=Lucida Sans Regular
serif.bold.3=Times New Roman WT TC
serif.bold.4=EUDC

serif.bolditalic.0=Times New Roman Bold Italic
serif.bolditalic.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
serif.bolditalic.2=Lucida Sans Oblique
serif.bolditalic.3=Times New Roman WT TC
serif.bolditalic.4=EUDC

sansserif.0=Arial
sansserif.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
sansserif.2=Lucida Sans Regular
sansserif.3=Arial Unicode MS
sansserif.4=EUDC

sansserif.italic.0=Arial Italic
sansserif.italic.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
sansserif.italic.2=Lucida Sans Oblique
sansserif.italic.3=Arial Unicode MS
sansserif.italic.4=EUDC

sansserif.bold.0=Arial Bold
sansserif.bold.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
sansserif.bold.2=Lucida Sans Regular
```

```

sansserif.bold.3=Arial Unicode MS
sansserif.bold.4=EUDC

sansserif.bolditalic.0=Arial Bold Italic
sansserif.bolditalic.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
sansserif.bolditalic.2=Lucida Sans Oblique
sansserif.bolditalic.3=Arial Unicode MS
sansserif.bolditalic.4=EUDC

monospaced.0=Courier New
monospaced.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
monospaced.2=Lucida Sans Typewriter Regular
monospaced.3=Lucida Sans Regular
monospaced.4=Monotype Sans Duospace WT TC
monospaced.5=EUDC

monospaced.italic.0=Courier New Italic
monospaced.italic.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
monospaced.italic.2=Lucida Sans Typewriter Oblique
monospaced.italic.3=Lucida Sans Oblique
monospaced.italic.4=Monotype Sans Duospace WT TC
monospaced.italic.5=EUDC

monospaced.bold.0=Courier New Bold
monospaced.bold.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
monospaced.bold.2=Lucida Sans Typewriter Regular
monospaced.bold.3=Lucida Sans Regular
monospaced.bold.4=Monotype Sans Duospace WT TC
monospaced.bold.5=EUDC

monospaced.bolditalic.0=Courier New Bold Italic
monospaced.bolditalic.1=Monotype Sans WT TC,CHINESEBIG5_CHARSET
monospaced.bolditalic.2=Lucida Sans Typewriter Oblique
monospaced.bolditalic.3=Lucida Sans Oblique
monospaced.bolditalic.4=Monotype Sans Duospace WT TC
monospaced.bolditalic.5=EUDC

#
# Define dialog and dialoginput to match the above definitions
#
alias.dialog=sansserif
alias.dialoginput=monospaced

#
# Exclusion Ranges
#
#exclusion.dialog.0=0500-20ab,20ad-ffff
#exclusion.dialoginput.1=0500-20ab,20ad-ffff
#exclusion.serif.0=0500-20ab,20ad-ffff
#exclusion.sansserif.0=0500-20ab,20ad-ffff
#exclusion.monospaced.1=0500-20ab,20ad-ffff

#
# To enable helvetica, courier and timesroman as real fonts
# remove aliases, and define them as real terminal fonts.
# (i.e. timesroman.0=TimesRoman)
#
# timesroman.0=TimesRoman
alias.timesroman=serif
# helvetica.0=Helvetica
alias.helvetica=sansserif
# courier.0=Courier
alias.courier=monospaced

# font filenames

filename.Arial=arial.ttf
filename.Arial_Bold=arialbd.ttf
filename.Arial_Bold_Italic=arialbi.ttf
filename.Arial_Italic=ariali.ttf
filename.Arial_Unicode_MS=arialuni.ttf
filename.Courier_New=cour.ttf
filename.Courier_New_Bold=courbd.ttf
filename.Courier_New_Bold_Italic=courbi.ttf
filename.Courier_New_Italic=couri.ttf
filename.EUDC=eudc.tte
filename.Lucida_Sans_Oblique=LucidaSansOblique.ttf
filename.Lucida_Sans_Regular=LucidaSansRegular.ttf
filename.Lucida_Sans_Typewriter_Oblique=LucidaTypewriterOblique.ttf
filename.Lucida_Sans_Typewriter_Regular=LucidaTypewriterRegular.ttf
filename.Monotype_Sans_Duospace_WT_T=mtsansdt.ttf
filename.Times_New_Roman=times.ttf
filename.Times_New_Roman_Bold=timesbd.ttf
filename.Times_New_Roman_Bold_Italic=timesbi.ttf
filename.Times_New_Roman_Italic=timesi.ttf
filename.Times_New_Roman_WT_TC=tnrwt_t.ttf
filename.Monotype_Sans_WT_TC=mtsans_t.ttf
filename.\u65b0\u7d30\u660e\u9ad4=mingliu.ttc
filename.MingLiU=mingliu.ttc
filename.PMingLiU=mingliu.ttc

# default char definition
default.char=02ff

# charset for text input

```

```
inputtextcharset=CHINESEBIG5_CHARSET

# font substitution
substitute.0=\u7d30\u660e\u9ad4=MingLiU
      substitute.1=\ub5b0\u7d30\u660e\u9ad4=PMingLiU
substitute.2=Arial Unicode MS=Times New Roman WT TC
substitute.3=Arial Unicode MS=Monotype Sans Duospace WT TC
substitute.4=Monotype Sans Duospace WT=Arial Unicode MS
substitute.5=Monotype Sans Duospace WT=Times New Roman WT TC
substitute.6=Times New Roman WT=Arial Unicode MS
substitute.7=Times New Roman WT=Monotype Sans Duospace WT TC
substitute.8=Times New Roman=Times New Roman
substitute.9=Courier New=Courier New
substitute.10=Arial=Arial
substitute.11=Times New Roman=Times New Roman
substitute.12=Courier New=Courier New
substitute.13=Arial=Arial
register.0=MingLiU
register.1=PMingLiU
===== end of fontproperties.zh_TW =====
```

5. Save the file and redisplay the help panels in IBM Spectrum Control GUI.

Cannot read Korean language logs from GUI

This problem occurs in Korean language environments.

Problem

You will not be able to view the log files from the IBM Spectrum® Control GUI because of a Db2® JDBC driver problem.

Action

You will have to read the files using another text viewer like **Notepad**.

Resolving problems

If you experience an error with IBM Spectrum® Control, you can use problem determination to determine why an error occurred, and explain how to resolve the problem.

Overview

The diagnostic tips that you can perform vary from including how to view and package log files, how to use network connectivity tools, and how to identify known problems and resolutions.

The first step in the problem determination process is to describe the problem completely.

Problem descriptions help you and IBM® Software Support know where to start to resolve the problem. Answer the following questions:

- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced?

Describe the symptoms of the problem

When you describe a problem, the most obvious question is "What is the problem?" By using the following questions, you can create a more descriptive picture of the problem:

- What are the error codes and messages?
- How does the system fail? For example, is it a loop, hang, crash, performance degradation, or incorrect result?
- What is the business impact of the problem?

Describe where the problem occurs

Determining where the problem originates is not always easy, but it is one of the most important steps in resolving a problem.

The following questions help you focus on where the problem occurs to isolate the problem:

- Is the problem specific to one operating system, or is it common across multiple operating systems?
- Is the environment and configuration supported?
- Is the application active locally on the database server or on a remote server?
- Is a gateway involved?
- Is the database on a local or remote computer?

Part of identifying where a problem originates is understanding the environment in which it exists.

Completely describe the problem environment, including the operating system and version, all corresponding software and versions, and hardware information. Confirm that the system is running within an environment that is a supported configuration.

Describe when the problem occurs

Develop a detailed timeline of events that led up to failure.

You can get a detailed timeline of events by starting at the time an error was reported. Use the available logs.

To develop a detailed timeline of events, answer these questions:

- Does the problem happen only at a certain time of day?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem occurred?
- Does the problem happen after an environment change such as upgrading or installing software or hardware?

These questions can help you set a frame of reference in which to investigate the problem.

Describe the conditions under which the problem occurs

Knowing which systems and applications are running at the time that a problem occurs is an important part of problem determination and troubleshooting. These questions about the operating system environment can help you identify the root cause of the problem:

- Does the problem always occur when the same task is being run?
- Must a certain sequence of events occur for the problem to occur?
- Do any other applications fail at the same time?

Can the problem be reproduced?

Problems that you can reproduce are often easier to debug and solve. However, problems that you can reproduce might have a disadvantage: If the problem has a significant business impact, you do not want it to recur.

If possible, re-create the problem in a test or development environment, which typically offers you more flexibility and control during your investigation.

Ask the following questions:

- Can the problem be re-created on a test system?
- Is the same type of problem encountered by multiple users or applications?
- Can the problem be re-created by running a single command, a set of commands, or a particular application?
- [IBM Spectrum Control tools](#)
Use the Service tool, the Repository copy tool and, the Tracing tool to help you diagnose errors with IBM Spectrum Control.
- [IBM Spectrum Control problem determination](#)
You can use this information to help you resolve IBM Spectrum Control problems.

IBM Spectrum Control tools

Use the Service tool, the Repository copy tool and, the Tracing tool to help you diagnose errors with IBM Spectrum® Control.

- [Repository copy tool](#)
You can use the Repository copy tool, **repocopy**, to export all the tables in the IBM Spectrum Control database repository for purposes of debugging problems.
- [Service tool overview](#)
The service tool collects information from all installed IBM Spectrum Control components. The tool detects the system configuration, collects the applicable information, and creates a compressed file that can be sent to IBM® Software Support.
- [Tracing servers and agents](#)
You can set trace levels for IBM Spectrum Control server and agent components. The tracing tool uses the IBM Logging Toolkit for C (CCLOG) for tracing server and agent data.

Repository copy tool

You can use the Repository copy tool, **repocopy**, to export all the tables in the IBM Spectrum® Control database repository for purposes of debugging problems.

You can send the exported data to IBM® Software Support to help debug problems.

Tip: If you want to export only performance data from the IBM Spectrum Control repository, you can create performance support packages. You can create performance support packages for storage systems or fabrics. For more information about exporting performance support packages, see [Exporting performance data for storage systems and fabrics](#).

- [Exporting repository data](#)
Use the Repository copy tool to export data from an existing repository into a text file.

Exporting repository data

Use the Repository copy tool to export data from an existing repository into a text file.

Procedure

To export repository data, follow these steps:

1. Go to the following default directory:

Windows operating systems:
c:\Program Files\IBM\TPC\data\server\tools

Linux® or AIX® operating systems:
/opt/IBM/TPC/data/server/tools

2. Issue the **repocopy** command:

Windows operating systems:
repocopy.bat

Linux or AIX operating systems:
repocopy

3. Select Export data from repository tables and click Next.
4. In the Options for Import/Export window, enter information in the following fields:

Directory for Export
Enter the directory where the comma-delimited file is saved.

Delimiter
Enter a delimiter for the delimited file format (a comma is the default).

Quote
Enter the symbol that contains string data (double quotation marks is the default).

IBM Spectrum® Control exports the data into the comma-delimited file that you specify, and places it in a file named *tablename.txt*. Click Next.

5. Select one of the following options and click Next.
 - Export by using DB2® native format.
 - Export by using text files (the preferred method).
6. Select one of the following options and click Next.
 - Export base tables (always export the base tables)
 - Export Performance Manager tables, if requested by IBM® Software Support
 - Export history tables, if requested by IBM Software Support.

The information that is detected in the server.config file is displayed in the Connection Properties window within the following fields:

- Database Types
- User name
- Password
- Driver Class
- Driver URL
- Database
- DB Creator
- Classpath

If you want to export data from a different database from the one listed in the server.config file, you can select the database from the Database Types list box. Manually enter the database information.

7. Click Finish.
8. Click Run.

Results

As you progress through the export process, messages are written to a progress log that is displayed. You can track the steps through the progress log. When the **repocopy** command is used with a remote database, the DB2 shared library is not available for loading the **libTSRMinsubd.so** file. You can ignore this message. Click OK and continue.

Service tool overview

The service tool collects information from all installed IBM Spectrum® Control components. The tool detects the system configuration, collects the applicable information, and creates a compressed file that can be sent to IBM® Software Support.

The service tool collects the following information:

- Host name
- IP address and configuration information
- Operating system and version. On the Windows operating system, a **msInfo.txt** report is also generated
- Java™ home, version, and class path
- Java Virtual Machine (JVM) implementation name and version
- Protocol statistics
- Internet Protocol network connections for IBM Spectrum Control, including listening ports
- Diagnostic information about the system and its services
- Listing of all library files, for example, server and library and agent and library
- HOSTS file
- IBM Spectrum Control version and license files

When the service tool is run on the system where the Data server or the Device server are installed, it also collects the following information:

- For the Data server, information about all of the remote and local graphical user interfaces (GUIs) that are associated with it
- For the Export server, information about this component configuration and logs
- For the Device server, Alert server, and web server, information about their profiles in IBM WebSphere® Application Server Liberty
- For the Data server, Device server, Alert server, web server, and data collector - all Javacore files, the most recent Java heap dump file, and the most recent snap trace file.

The service tool can collect Java core dump files for these IBM Spectrum Control components, but it *does not* by default. The values for the `javaCoreFiles`, `heapDumpFiles`, `snapFiles`, and `coreFiles` parameters in the `installation_dir/service/service.properties` file dictate which types of files and how many of each file type are collected by the service tool.

- All applied interim fixes
- Installation logs
- The contents of the `log` and `logs` directory, including subdirectories
- The contents of the `conf` and `config` directory
- Directory listing of the `lib` and `bin` directory
- The contents of the `log` and `conf` subdirectories of the `web` directory
- For the IBM Spectrum Control GUI, information about its profile in the embedded WebSphere Application Server Liberty
- Information from the `ipconfig /all` command on Windows operating systems
- Information from the `ipconfig -a` command on Linux® and AIX® operating systems
- Information from the `netstat -an` command on all operating systems

When the service tool is run on the system where the database repository is installed, it also collects the DB2® support information.

When the service tool runs on the Storage Resource agent computer, it collects the following information:

- All applied interim fixes
- Everything in the `config`, `log`, `nls`, `output`, and `service` directories, including subdirectories
- Everything in the `opt/IBM/CAP` directory on Linux and AIX operating systems
- Directory listing of the `ProgramData\Application Data\IBM\CAP` directory on Windows operating systems
- Directory listing of the `agent` directory
- Directory listing of the `bin` directory
- Listing of version numbers for the Storage Resource agent component
- Information from the `ipconfig /all` command on Windows operating systems
- Information from the `ifconfig -a` command on Linux and AIX operating systems
- Information from the `netstat -an` command on all operating systems

By default, the service data is collected in one of the following directories:

For Windows operating systems:

`installation_dir\service\data`

For Linux and AIX operating systems:

`installation_dir/service/data`

For more information about changing the default directory, see [How to customize the service tool](#).

You can run the service tool on IBM Spectrum Control regardless of whether you configured it on a single server or on multiple servers. The service tool automatically recognizes the installed components and collects service data about them. For more information about running the service tool for servers, see [Packaging log files from the command line and sending them to IBM Support](#).

- [Packaging log files from the command line and sending them to IBM Support](#)
To provide trace information to IBM Support about the performance of IBM Spectrum Control, you might be asked to package and send a set of log files. You can run the service tool from the command line to create the log package and then manually upload it to IBM Support.
- [Creating a compressed file for a Storage Resource agent](#)
Run the service tool on Storage Resource agents that were deployed by using the web-based GUI to create a compressed file that can be sent to IBM Software Support.
- [How to customize the service tool](#)
The default behavior of the service tool is to collect data about all IBM Spectrum Control components, but you can use the service tool to collect data about specific components. You can also use command-line parameters to specify a location to place the data that is collected, specify that the data is compressed, or to specify both.

Packaging log files from the command line and sending them to IBM Support

To provide trace information to IBM® Support about the performance of IBM Spectrum® Control, you might be asked to package and send a set of log files. You can run the service tool from the command line to create the log package and then manually upload it to IBM Support.

Before you begin

To package IBM Spectrum Control log files, you must have administrator authority on Windows operating systems or root authority on AIX® and Linux® operating systems. Before packaging or sending the log files, you must first open a support ticket with IBM Support. For information, see [Getting support](#).

In order to run the service tool when you log in by using a Windows domain user account, you must grant Db2® SYSADM authority to that Windows domain user account.

To run the service tool when you log in by using a domain user account, choose one of the following methods:

- Right-click the `service.bat` file and select Run as administrator.
Or
- Click Start > All Programs > Open Administrative Tools > Local Security Policy. For information about how to open Administrative Tools, see [Accessing administration tools](#).

- On the Local Security Policy window, disable User Account Control: Run all administrators in Admin Approval Mode.
- Restart your computer.

About this task

The compressed file contains data about the following IBM Spectrum Control components: Alert server, Data server, Device server, Web server, Storage Resource agent, DB2®, CLI, and installation.

Tip: You can customize the service tool to collect data about specific IBM Spectrum Control components. For more information, see [How to customize the service tool](#).

Procedure

To run the service tool for all components, follow these steps:

1. Log on to the system where IBM Spectrum Control is installed.
2. Go to the following directory:

Windows operating systems:
installation_dir\service

Linux or AIX operating systems:
installation_dir/service/

3. Run the following program:

Windows operating systems:
 service.bat

Linux or AIX operating systems:
 service.sh

A compressed file, SCServiceFiles_all.zip, is created in the following directory:

Windows operating systems:
installation_dir\service\data

Linux or AIX operating systems:
installation_dir/service/data/

4. Upload the log package by using one of the following methods:

If . . .	Then . . .
You are a US healthcare client	Blue Diamond data upload instructions .
The log file is < 200 MiB	Upload and attach the file to the case .
The log file is > 200 MiB	Enhanced Customer Data Repository (ECuRep) - Send data (FTP) .

Related tasks

- [Creating a compressed file for a Storage Resource agent](#)

Related reference

- [How to customize the service tool](#)

Creating a compressed file for a Storage Resource agent

Run the service tool on Storage Resource agents that were deployed by using the web-based GUI to create a compressed file that can be sent to IBM® Software Support.

Before you begin

You must have administrator authority on Windows operating systems or root authority on AIX® and Linux® operating systems.

Procedure

To run the service tool on the Storage Resource agents, complete the following steps:

1. In the menu bar, go to Servers, > Servers.
2. Right-click the server where the Storage Resource agent is deployed, and select Logs, > Collect Agent Logs.

Results

The following compressed files are created:

Windows operating systems:
 C:\Program Files\IBM\TPC\data\log\SRATraces*SRA_computer_name*\SCServiceInfo.zip

AIX and Linux operating systems:
 /opt/IBM/TPC/data/log/SRATraces/*SRA_computer_name*/SCServiceInfo.zip

where `SRA_computer_name` represents the name of the computer on which the Storage Resource agent is located. If the compressed file cannot be created, a message indicates that the job was unsuccessful.

For more information about the error, see the server log file or the services script information file. The files are in one of the following default directories:

Server log file

This file is on the system where IBM Spectrum® Control is installed:

Windows operating systems:

`c:\Program Files\IBM\TPC\data\log`

AIX or Linux operating systems:

`/opt/IBM/TPC/data/log`

Services script information file

This file is on the computer on which the Storage Resource agent is installed:

Windows operating systems:

`c:\Program Files\IBM\TPC\SRA_computer_name\services\SCServiceInfo.log`

AIX or Linux operating systems:

`/opt/IBM/TPC/SRA_computer_name/services/SCServiceInfo.log`

For more information about customizing the data that is collected by the service tool, see [How to customize the service tool](#).

Related tasks

- [Packaging log files from the command line and sending them to IBM Support](#)

How to customize the service tool

The default behavior of the service tool is to collect data about all IBM Spectrum® Control components, but you can use the service tool to collect data about specific components. You can also use command-line parameters to specify a location to place the data that is collected, specify that the data is compressed, or to specify both.

Specifying help and output command-line parameters

To obtain information about the service tool usage, use the `-help` command-line parameter.

To specify the data that is collected by the service tool, use the following command-line parameters when you run service tool:

`-output directory_path`

Places the files that contain the data that was collected in a directory that you specify. If you specify a directory that does not exist on your system, that directory is created. If you do not use the `-output directory_path` parameter, the files are placed in the default directory:

Windows operating systems

`installation_dir\service\data`

Linux® or AIX® operating systems

`installation_dir/service/data`

Restriction: If you specify a directory, the directory path cannot contain spaces. This restriction refers to the `-output` option.

`-pmr`

The number of the PMR to which the support information is related. Use the following format to enter the number: `nnnnnn,nnn,nnn` where `n` represents a number, such as `12345,123,123`. If you enter a value for the `-pmr` parameter, the service tool automatically uses FTP to upload the service information to IBM®. If the service tool cannot upload the service information to IBM from the IBM Spectrum Control server, the generated file must be uploaded manually. The `-pmr` parameter is ignored if the `-nozip` parameter is used.

`-nozip`

When you use this parameter, the compressed collected data archives are no longer created. The service tool creates separate directories for each component for which data was collected. You can then create compressed archives for the collected files. In this way, you can control the size and content of each compressed file. To specify a directory other than the default directory, use the `-output directory_path` parameter.

Tip: You can specify more than one parameter, for example, `C:\Program Files\IBM\TPC\service>service -install -nozip`.

Specifying Javacore, Java heap dump, snap trace, and Java core dump file collection

By default, the service tool collects all Java™ core files, the most recent Java heap dump file, and the most recent snap trace file for the Data server, Device server, Alert server, web server, and data collector. The service tool can collect Java core dump files for these IBM Spectrum Control components, but it *does not* by default. You can edit the following parameters in the `installation_dir/service/service.properties` file in order to control which types of files and how many of each file type are collected by the service tool:

`coreFiles=0`

Specifies how many Java core dump files are collected. The default value is 0 because Java core dump files are generally very large. The valid values for this parameter are 0 or a positive integer. Any other value results in the default behavior.

`javaCoreFiles=-1`

Specifies how many Javacore txt files are collected. The default value of -1 means that all Javacore txt files are collected. The valid values for this parameter are 0 or a positive integer. Any other value results in the default behavior.

`heapDumpFiles=1`

Specifies how many Java heap dump files are collected. The default is to collect the most recent Java heap dump file. The valid values for this parameter are 0 or a positive integer. Any other value results in the default behavior.

`snapFiles=1`

Specifies how many snap trace files are collected. The default is to collect the most recent snap trace file. The valid values for this parameter are 0 or a positive integer. Any other value results in the default behavior..

Collecting data for specific IBM Spectrum Control components

You can use the service tool to collect data about specific IBM Spectrum Control components.

Use the following parameters to specify the components:

-all
All components. The default behavior is to collect data about all components.

-install
Installation component files.

-data
Data server component files.

-device
Device server component files.

-datacollector
Data collector component files.

-alert
Alert server component files.

-export
Export server component files.

-sra
Storage Resource agent component files.

-db
Db2® files.

-cli
Command-line interface files.

-gui
Collects data about GUI files.

If you collect information about a particular component, and you do not specify the -nozip parameter, you can identify the contents of a compressed service file from its name. For example, if you specify the -db -gui parameters but did not specify the -nozip parameter, a file named SCSERVICEFiles_db_gui.zip is created.

Warning: An existing compressed file is overwritten when another file of the same name is created. For example, if you run **C:\Program Files\IBM\TPC\service>service -db -gui**, a file named SCSERVICEFiles_db_gui.zip is created. If you rerun the tool with the same component options, a new file named SCSERVICEFiles_db_gui.zip is created. This new file overwrites the previously created file unless you specify the -nozip parameter, or use the -output parameter to specify a different path. If you specify the -nozip parameter, a directory is created for the components that you specified in the parameter. If you did not specify a specific component, data is collected for all installed components, and the data is placed in files in the following directory:

Windows operating systems:
installation_dir\service\data

Linux or AIX operating systems:
installation_dir/service/data

Files for particular components are then placed in a directory that corresponds to that component. Common files, such as license.txt, are placed in the following directory:

Windows operating systems:
installation_dir\service\data

Linux or AIX operating systems:
installation_dir/service/data

Restriction: You cannot specify a specific component as a command-line parameter when you specify the -all parameter. Also, when you specify a component that is not installed on the computer, the service tool displays an error message.

Tracing servers and agents

You can set trace levels for IBM Spectrum® Control server and agent components. The tracing tool uses the IBM® Logging Toolkit for C (CCLOG) for tracing server and agent data.

Setting server trace levels in the GUI

If you have Administrator permissions, you can turn on or off the trace recording for the individual services on each component server. By enabling or disabling the trace recording before you create logs, you can provide IBM Support with more information on specific component servers that might be causing the problem.

To enable or disable the trace recording for a component servers, complete the following steps:

1. In the menu bar, go to Home > System Management.
2. Click Component Servers in the Components section. In the Component Servers pane, you can view performance information for each component server. The current trace recording for each component server is highlighted.
3. Click On or Off to enable or disable tracing for each of the services of a component server.

Tracing a Storage Resource agent

To trace a Storage Resource agent, follow these steps:

1. In the menu bar, go to Servers, > Servers.
2. In the Servers pane, right-click a server and click View Properties.
3. Click the Agent tab and click Edit.
4. Specify the following values:

Agent Trace

Click Enabled.

Trace Level

The trace level can be one of the following values:

Minimum

The minimum debugging level. This option is the default and is always turned on. This option includes trace information for key locations in the code that includes first failure data capture. This setting does not capture any output data.

Medium

The medium debugging level. This setting does not capture any output data.

Maximum

The maximum debugging level. This option includes entry and exit information for every method in addition to all **DEBUG_MIN** statements. This option can impact performance.

Number of Trace Files

The maximum number of files that are used before the Data server starts reusing the tracing files. The default is eight files.

Trace File Size (MiB)

The maximum file size for the tracing files. The default is one MiB.

5. Click Save and then click Close.

Trace log file locations

The Java™ tracing information is stored in the locations that are shown in [Table 1](#) and [Table 2](#).

Table 1. Java trace log files

Class	Java trace log file
Storage Resource agent	<i>installation_directory/agent/log</i> <i>installation_directory/agent/log/</i> <i>name_of_server_SRA_communicates_with</i>
Server	<i>installation_directory/data/log/ServerTrace.log</i>

The Native tracing information is stored in the locations that are shown in [Table 2](#).

Table 2. Native trace log files

Class	Native trace log file
Server	<i>installation_directory/data/log/traceTPCDNativeServer_ yyyymmddhhmmssLOGx.log</i>
UnixStopServer	<i>installation_directory/data/log/traceTPCDNativeStopServer_ yyyymmddhhmmssLOGx.log</i>
ImportExport	<i>installation_directory/data/log/traceTPCDNativeImportExport_ yyyymmddhhmmssLOGx.log</i>

Messages are logged before the creation of a user-defined installation directory. Because the log directory is not initially available, installation messages are written to standard error log files. The standard error log is redirected to a temporary file. The temporary message log file is copied to the appropriate log directory, and the standard error log file is redirected to reflect the new message log file location

When an executable file is run, it creates the tracing files that are shown in [Table 3](#).

Table 3. Trace files for executable routines

Executable file	Platform	Trace log file location, file name, and description
udbexec.exe	Windows operating systems	<i>\hostname\traceTSRMNativeUdbexec_PID.log</i> <i>installation_directory\data\ca\</i> <i>traceTSRMNativeUdbexec_PID_ yyyymmddhhmmss.log</i>
ExecSvc.exe	Windows operating systems	<i>temp\traceTSRMNativeExecSvc.log</i> This executable routine runs only on the remote computer during a Windows push installation.
udbexec.exe ¹ or udbexec ²	All operating systems	<i>installation_directory/data/ca/</i> <i>traceTSRMNativeudbexec_pid_ yyyymmddhhmmss.log</i>
DbMonitor.exe ¹ or DbMonitor ²	All operating systems	This executable routine is run on Windows operating systems for all database monitoring activities. This executable routine generates trace files that are named based on the name of the active job. For example, the default database scan generates a trace file similar to the following file: <i>TPCUser.Default Db Scan.0001_dbmonitor.00009163.trace</i>

Notes:

- ¹ Use in Windows operating systems.
- ² Use in Linux® or AIX® operating systems.

Configuration files for tracing

There are two configuration files added for Java tracing: ServerTraceLog.config and AgentTraceLog.config. The nativelog.config file is used for native code tracing.

Table 4. Trace configuration files

File	Directory
Native code trace configuration file	<i>installation_directory/data/config/nativelog.config</i>
Java tracing configuration file	<i>installation_directory/data/config/ServerTraceLog.config</i>

nativelog.config file

During the installation or a maintenance upgrade of IBM Spectrum Control, a file that is named `nativelog.config` is created in the `installation_directory/data/config` directory.

This file is used to specify the default levels of tracing for the IBM Spectrum Control processes (for example, the server and the agent). Depending on which platform IBM Spectrum Control is installed, some of the items that are listed in this file might vary.

In particular, Windows installations do not have the `#StopAgent` section that Linux and AIX platforms have. The reason for the difference is because the agent service is started differently on these platforms. Settings in these files are separated into sections that represent IBM Spectrum Control processes. The trace settings per process are listed underneath the process they represent.

The definitions for the entries in the `nativelog.config` file are as follows:

```
# Agent
    A IBM Spectrum Control process. In this case, the agent process is specified.
agt.level=DEBUG_MIN
    The default trace level for the agent process.
agt.maxFiles=3
    The maximum number of trace files to be created or used.
agt.maxFileSize=20480000
    The maximum size of each trace file (in bytes).
```

The level of trace that is specified in this file is also applied to specific child processes run from their parent process. For example, the processes `udbexec.exe` and `MiniProbe.exe` can use the settings from the Agent section as their default trace level because these processes are run by the Agent process.

You can edit the contents of this file if you want to change the default level of tracing for a particular process. However, these changes do not take effect until that process is stopped and restarted again. For example, if you want to change the tracing level of the server from `DEBUG_MIN` to `DEBUG_MID`, you can modify the `nativelog.config` file. Restart the server to effect the change.

By modifying this file, you can ensure that each time the server is stopped and restarted, the specified trace level is applied.

Here is an example of the `nativelog.config` file for Windows operating systems:

```
# level: DEBUG_MIN || DEBUG_MID || DEBUG_MAX
# maxFiles: >= 2 && <= 100
# maxFileSize: >= 128000 && <= 102400000

# Agent
agt.level=DEBUG_MIN
agt.maxFiles=3
agt.maxFileSize=20480000

# Server
srv.level=DEBUG_MIN
srv.maxFiles=3
srv.maxFileSize=20480000

# ImportExport
impexp.level=DEBUG_MIN
impexp.maxFiles=3
impexp.maxFileSize=20480000
```

Here is an example of the `nativelog.config` file for Linux and AIX operating systems:

```
# level: DEBUG_MIN || DEBUG_MID || DEBUG_MAX
# maxFiles: >= 2 && <= 100
# maxFileSize: >= 128000 && <= 102400000

# Agent
agt.level=DEBUG_MIN
agt.maxFiles=3
agt.maxFileSize=20480000

# ImportExport
impexp.level=DEBUG_MIN
impexp.maxFiles=3
impexp.maxFileSize=20480000

# StopAgent
stpagt.level=DEBUG_MIN
stpagt.maxFiles=3
stpagt.maxFileSize=20480000
```

AgentTraceLog.config file

The definitions for the entries in the `AgentTraceLog.config` file are as follows:

```
ITSRM.logger.trace.Agent.logging=true
    The state of tracing for the agent. In this case, tracing is enabled.
ITSRM.logger.trace.Agent.level=DEBUG_MAX
    The trace level for the agent. DEBUG_MAX is the only level that produces trace information.
ITSRM.handler.file.maxFiles=8
    The maximum number of trace files to be created or used.
ITSRM.handler.file.maxFileSize=1024
    The maximum size of each trace file (in bytes).
```

ServerTraceLog.config file

The definitions for the entries in the ServerTraceLog.config file are as follows:

ITSRM.logger.trace.Server.logging=true

The state of tracing for the Server process. In this case, tracing is enabled.

ITSRM.logger.trace.Server.level=DEBUG_MAX

The trace level for the Server process. **DEBUG_MAX** is the only level that produces trace information.

ITSRM.handler.file.maxFiles=5

The maximum number of trace files to be created or used.

ITSRM.handler.file.maxFileSize=20480

The maximum size of each trace file (in bytes).

ITSRM.logger.trace.TivoliSRM-GUI.logging=false

The state of tracing for the stand-alone GUI process. In this case, tracing is disabled.

ITSRM.logger.trace.TivoliSRM-GUI.level=DEBUG_MAX

The trace level for the stand-alone GUI process.

ITSRM.logger.trace.TivoliSRM-CIMOM.logging=false

The state of tracing for the CIMOM process. In this case tracing, is disabled.

ITSRM.logger.trace.TivoliSRM-CIMOM.level=DEBUG_MAX

The trace level for the CIMOM process.

ITSRM.logger.trace.TivoliSRM-Agent.logging=false

The state of tracing for the Agent process. In this case, tracing is disabled.

ITSRM.logger.trace.TivoliSRM-Agent.level=DEBUG_MAX

The trace level for the Agent process.

ITSRM.logger.trace.Scheduler.logging=false

The state of tracing for the Scheduler process. In this case, tracing is disabled.

ITSRM.logger.trace.Scheduler.level=DEBUG_MAX

The trace level for the Scheduler process.

IBM Spectrum Control problem determination

You can use this information to help you resolve IBM Spectrum® Control problems.

- [Configuration files](#)
Use the parameters in IBM Spectrum Control configuration files to help resolve problems.
- [Log files](#)
When you have a problem, you can check several product log files.
- [Audit logs](#)
The audit logs provide an audit of most activities. The logs include the identification of users that initiate actions in the product.
- [Diagnosing IBM Spectrum Control problems](#)
You can use the log files, audit files, configuration files, and what steps to take, to diagnose a IBM Spectrum Control problem.

Configuration files

Use the parameters in IBM Spectrum® Control configuration files to help resolve problems.

The parameters in the configuration files are case-sensitive.

The default *installation_dir* for IBM Spectrum Control installations is as follows:

Windows operating systems:

c:\Program Files\IBM\TPC

Linux® or AIX® operating systems:

/opt/IBM/TPC

The default file locations for the configuration files for IBM Spectrum Control are as follows:

IBM Spectrum Control:

Windows operating systems:

installation_dir\config

Linux or AIX operating systems:

installation_dir/config

Data server:

Windows operating systems:

installation_dir\data\config

Linux or AIX operating systems:

installation_dir/data/config

Device server:

Windows operating systems:

installation_dir\device\conf

Linux or AIX operating systems:

installation_dir/device/conf

Storage Resource agent:

Windows operating systems:
installation_dir\agent\config\agent.config

Linux or AIX operating systems: operating systems:
installation_dir/agent/config/agent.config

Restriction: On Windows installations, if you installed IBM Spectrum Control by using a domain user account, you must disable User Account Control to edit the configuration files.

- [server.config file](#)
The following information lists the parameters that are set in the server.config file. These parameters include controller, logging, repository, and service.
- [scheduler.config file](#)
The following information lists the parameters that are set in the scheduler.config file. These parameters include concurrency parameters and jobs parameters.
- [TPCD.config file](#)
The list of parameters that are set in the TPCD.config file include server parameters and GUI parameters.
- [agent.config file](#)
The agent.config file contains configuration parameters for the Storage Resource agent. These parameters are set when the Storage Resource agent is installed. The parameters can also be changed manually by editing the file.

server.config file

The following information lists the parameters that are set in the server.config file. These parameters include controller, logging, repository, and service.

Controller parameters

name
The Data Manager server name is the name of the host computer.

port
The port on which the server listens for requests. The default is 9549.

maxConnections
The maximum number of concurrent sockets that the server opens. The default is 500.

routerThreads
The number of threads that redirect incoming requests to the appropriate service provider. The default is 1.

serviceThreads
The number of threads to allocate for the internal service provider of the server. The default is 2.

agentErrorLimit
The number of consecutive attempts to reach an agent before the agent is displayed as DOWN. The default is 3. When an agent is in this state, no attempts to connect are made until either the agent contacts the server or the agent status is manually changed to UP.

adminGroup
The name of the group a user must be a member of to perform administrative functions from the GUI, the default is **adm**.

commEncrypted
The switch that secures communication between the Server or Agent and the Server/GUI by encrypting the DataStream.

- **0** = Off. Do not encrypt the DataStream.
- **1** = On. Encrypt the DataStream.

FileSystemScan NFSTimeout
Determines the numbers of seconds that a Storage Resource agent waits for a status system call on a Network File System (NFS) before it times out.

hostAlias
This parameter is displayed if the HOST_ALIAS is not specific and represents the name of the server. The value for this parameter is used when multiple computers have the same name or the name cannot be determined.

Logging parameters

logsKept
The number of server logs to keep. The default is 5.

messagesPerLog
The maximum number of messages in a log. When this number is reached the log is closed and a new log is created. The default is 100,000.

Repository parameters

driver
The name of the JDBC driver to use, normally:

- Db2®: **COM.ibm.db2.jdbc.app.DB2Driver**

url
The URL used to connect to the database, normally:

- Db2: **jdbc:db2:database_name**

user

The user name that IBM Spectrum® Control uses to connect to the repository.

connectionPool

The number of database connections in a pool of reusable open connections. The default is 10.

Service parameters

name

Repeating section that indicates the service providers that are required to start.

The REQUIRED parameters are as follows:

- `TStorm.server.svp.GuiSvp`
- `TStorm.server.svp.AgentSvp`
- `scheduler.Scheduler`

scheduler.config file

The following information lists the parameters that are set in the scheduler.config file. These parameters include concurrency parameters and jobs parameters.

Concurrency parameters

maxSubmitThreads

The maximum number of threads to create that handle the submission of jobs. The default is 3.

maxCompleteThreads

The maximum number of threads to create to handle job completions. Initially creates a pool of half the number of threads specified that can grow to the maximum. The default is 3.

Jobs parameters

minutesAdvanced

The number of minutes in advance of scheduled time to begin the scheduling process. The default is 1. Use this option to allow for the processor time that is involved in scheduling a job so that the job starts close to the scheduled time.

delayLimitMinutes

Number of minutes after scheduled start time that the Scheduler continues to attempt to start a job for a selected resource, so that resource state is not in a down state or, connection status is not in a failed state. The default is 120.

Location of the scheduler.config file

The scheduler.config file is in the following directories:

Windows operating systems:

installation_dir\data\config

Linux® or AIX® operating systems:

installation_dir/data/config

TPCD.config file

The list of parameters that are set in the TPCD.config file include server parameters and GUI parameters.

The following list describes the server parameters:

threadPoolSize

Number of initial threads to create for handling requests. The default is 3.

abbreviatedProbe

Only SCSI commands are sent to disk drives for inquiry and disk capacity information. The default is 1.

maxThreads

Set the maximum number of threads for handling requests. The default is 8.

pingReceiveTimeout

Number of seconds to wait before it indicates that a ping failed. The default is 10.

skipAutoFS

Set to 1 if you want to skip the **automount** process during discovery on the Oracle Solaris Storage Resource agent. By default, discovery always processes **automount** on all Oracle Solaris Storage Resource agents managed by the Data server.

saveNonRoot

Set to 1 if you want to monitor non root exports. The default is 0.

If you do not set this parameter, the export paths that are not at the root of the file system are discarded. If the NAS server has only non root exports accessible to the agent, it will not be added. Restart the Data server for this setting to take effect.

batchPartitionWaitRetryCount

Specify the number of times that the Storage Resource agent tries to get a report partition before the Storage Resource agent fails with an error.

Large batch reports are generated in partitions. The partitions are placed on the IBM Spectrum® Control server, and the Storage Resource agent gets them from the server when the batch report is created.

The following list describes the GUI parameters:

threadPoolSize

Number of initial threads to create for handling user interface requests. The default is 3.

maxThreads

Set the maximum number of threads for handling user interface requests. The default is 10.

reportRowLimit

Maximum number of rows that are sent at a time to the user interface. If this number is exceeded, a More button is displayed over the table, along with a warning message. The default is 5000.

keepCachedReport

Number of minutes to retain incomplete reports in the tmp directory for the server. The default is 120.

agent.config file

The agent.config file contains configuration parameters for the Storage Resource agent. These parameters are set when the Storage Resource agent is installed. The parameters can also be changed manually by editing the file.

The following list contains the parameters for the agent.config file.

Servename

Fully qualified host name of the system on which the Data server is installed.

Portnumber

Port on which the Data server listens for communications from the Storage Resource agent. By default, the port is set to 9549.

IPAddress

IP address of the server on which the Data server is installed.

Log files

When you have a problem, you can check several product log files.

- **[Installation log files for IBM Spectrum Control](#)**

You can use the installation log files, that contain all the installation messages and traces that are generated during the installation process, to debug installation problems.

Related reference

- [Default locations of log files](#)

Installation log files for IBM Spectrum Control

You can use the installation log files, that contain all the installation messages and traces that are generated during the installation process, to debug installation problems.

The lax*-out.txt and lax*-err.txt files

The lax*-out.txt and lax*-err.txt files are Installation Anywhere log files that contain all of the traces that are generated by the IBM Spectrum® Control installer.

These log files are useful for debugging any issues that occurred before you clicked Install on the installer. You can find them in the following locations:

Windows operating systems (assuming you are installing as Administrator):

C:\Users\Administrator\AppData\Local\Temp\2

Linux® or AIX® operating systems:

/tmp

The msgTPCInstall.log and traceTPCInstall.log files

The msgTPCInstall.log and traceTPCInstall.log files are Message and Trace Installation logs that contain all of the installation messages and traces, after you clicked Install on the installer. You can find them in the following locations:

Windows operating systems:

TPC_installation_directory\TPC\logs

Linux or AIX operating systems:

TPC_installation_directory/logs

The agent*.log file

The agent*.log file is the storage resource agent installation log file. You can find it in the following locations:

Windows operating systems:

TPC_installation_directory\logs\install\sra

Linux or AIX operating systems:
`TPC_installation_directory/logs/install/sra`

The `componentInstall.log` and `componentInstallIS.log` files

The `componentInstall.log` and `componentInstallIS.log` files are the main IBM Spectrum Control log files for installation. `component` can be any of the following components:

- `ca`
- `cli`
- `data`
- `dbSchema`
- `device`
- `gui`
- `agent` (Storage Resource agent)

You can find the `componentInstall.log` and `componentInstallIS.log` files in the following location:

Windows operating systems:
`TPC_installation_directory\log\component\install`
Linux or AIX operating systems:
`TPC_installation_directory/log/component/install`

Audit logs

The audit logs provide an audit of most activities. The logs include the identification of users that initiate actions in the product.

These are the audit logs:

- Web server, `traceWebServer_<number>.log`, which is initiated by the GUI, where `<number>` is sequential from 0 to 9 with 0 used for the most recent log file.
- Data server, `AuditTrace.log`, which is initiated by using the GUI.
- Device server, `auditTPCDeviceServer.log`, which is initiated by the CLI.

Audit logging is performed as follows:

- GUI commands are logged in the Data and Web server audit log.
- CLI commands are logged in the Device server audit log.

The communication between the following services are not logged:

- Data server to Device server
- Data server to agents
- Device server to agents

The audit logs are in the following directories:

Web server audit log

`installation_dir\web\log\traceWebServer_<number>.log`, where `<number>` is sequential from 0 to 9 with 0 used for the most recent log file.
The Web server audit log includes the following information:

- Timestamp
- User ID
- Service and operation performed
- Key input and output parameters

For example:

```
2020-11-06 10:18:52.710-0700 [Default_Executor-thread-47] [TPCAuthenticationService.authenticate] [INFO] User db2inst1 is authenticated with password
```

```
2020-11-09 05:47:11.638-0700 [pool-9-thread-7] [JobManagerServiceImpl.actionRunNow] [INFO] argEntities: [SANEntityInfo [parent=null, entity.id=33027, entity.type=DISKSYSTEM, entityName=FLYER 7825410, message=null, osType=XIV]] jobType: 16
```

Data server audit log:

`installation_dir\data\log\AuditTrace.log`

The Data server audit log includes the following information:

- Timestamp
- User ID
- Service and operation performed

For example:

```
2016-04-19 05:59:00.954-0700 Username=TPCUser
Running Ping com.tiv.itsrm.TStorm.server.svp.JobHndlr
handle AgentRequestHandler52016-04-19 05:59:00.978-0700
Username=TPCUser
Running com.tiv.itsrm.TStorm.server.svp.Ping
com.tiv.itsrm.TStorm.server.svp.JobHndlr
handle com.tiv.itsrm.TStorm.server.svp.JobHndlr
2016-04-19 06:03:29.404-0700 Username=TPCUser
Create. Definition type=PROBE.
Name='Probe_linux_vcloud211.storage.tucson.abc.com'
com.tiv.itsrm.TStorm.server.svp.GuiProbeHndlr save()
```



```
GuiRequestHandler9
2016-04-19 06:09:01.029-0700 Username=TPCUser
Running Ping com.tivoli.itsrm.TStorm.server.svp.JobHndlr
handle AgentRequestHandler72016-04-19 06:09:01.031-0700
Username=TPCUser
Running com.tiv.itsrm.TStorm.server.svp.Ping
com.tiv.itsrm.TStorm.server.svp.JobHndlr
handle com.tiv.itsrm.TStorm.server.svp.JobHndlr
2016-04-19 06:23:00.159-0700 Username=TPCUser
```

This audit log exists immediately after the product is installed in your environment.

Device server audit log:

installation_dir\device\log\auditTPCDeviceServer.log

The Device server audit log includes the following information:

- Timestamp
- User ID
- IP address of client
- Service and operation performed
- Key input and output parameters

You can only generate the auditTPCDeviceServer.log file with the `tpctool`.

For example:

```
tpctool.bat lsvol -dev subsystem_ID+0 -l -url localhost:9550
-user <dbadmin> -pwd <db2admin password>
```

```
2016-04-06 00:07:22.004-0700 BTACS0047I
DiskManagerService.getStorageVolumesForSubsystem
performed by db2admin at 127.0.0.1.
Input parameters: [volume_ID x+0, {DETECTABLE=1}], output parameters:
[volume_ID x+9+subsystem_ID+0,
volume_ID x+9+subsystem_ID+0,volume_id+9+subsystem_ID+0,
volume_id x+9+subsystem_ID+0, 6
```

Tip: Not all commands are logged into the auditTPCDeviceServer.log file. For example:

```
tpctool.bat catdscfg -url localhost:9550 -user user name -pwd user password
tpctool.bat setdscfg -url localhost:9550 -user user name -pwd user password
-property SnmpRetryCount -context DeviceServer 999999999
```

This audit log does not exist immediately after you install the product; it is generated when you run your initial CLI command.

For more information, see [tpctool](#).

Diagnosing IBM Spectrum Control problems

You can use the log files, audit files, configuration files, and what steps to take, to diagnose a IBM Spectrum® Control problem.

When a problem occurs with IBM Spectrum Control, collect pertinent information to help diagnose the problem. The following topics help you gather the information that you need.

For specific troubleshooting topics, see the IBM Spectrum Control Information Center.

General information requirements for problem diagnosis

When a problem occurs with IBM Spectrum Control, gather the following information:

- An exact description of the problem.
- The function in use at the time that the problem occurs.
- The sequence of steps that resulted in the problem.
- The expected results from the failing step.
- Any error messages that you see.
- The date and time when the problem occurred.
- The log files that are collected with the service utility.
- The last time that inventory collection was run, which indicates that the repository is in synchronization with the real configuration.
- Whether the error is repeatable or it occurs intermittently.
- The answers to the following connectivity questions help you determine whether there is a communication issue between IBM Spectrum Control and the CIM Agent (or SMI-S provider) server:
 - Is there any firewall that is enabled interfering with the communication between the IBM Spectrum Control components and the SMI-S provider?
 - Is it possible to ping the SMI-S provider?
 - Is it possible to telnet to the SMI-S provider?
 - Is it possible to contact the SMI-S provider with a CIM browser?
 - Are the SMI-S provider in the local subnet or is a Directory Agent (DA) in another subnet used for discovery?
 - Is the DA on the remote subnet registered with IBM Spectrum Control?
 - Was the SMI-S provider Service Location Protocol (SLP) registration successful? Use the **slptool findsrvs** command.

When you schedule data collection, the actual start time of the job starts 1 minute earlier, which is the default, than the scheduled time. So, if you schedule data collection to run at 1:30 PM, the job actually starts at 1:29 PM. The 1:29 PM time stamp is displayed in the GUI for that job and in the job log files. The action takes into account the system processor time that is involved in scheduling a job so that the job actually starts close to the scheduled time.

- [Discovery](#)
Use this information to troubleshoot discovery problems.
- [Discovery and probe completion codes](#)
Use this information to interpret status and completion codes returned in messages from discovery and probe jobs.
- [Monitoring service](#)
The following information shows you how to troubleshoot monitoring service problems.
- [Performance monitoring](#)
The following information shows you how to troubleshoot performance monitoring problems.
- [Configuration History](#)
Diagnose problems with the Configuration History view.
- [SMI-S fabric probe](#)
The following information shows you how to diagnose problems with the SMI-S probes of fabrics.
- [VMware ESX](#)
The following information shows you how to diagnose problems with VMware ESX servers.
- [FlashCopy](#)
IBM Spectrum Control FlashCopy® support labels volumes as having a FlashCopy property of *source*, *target*, or *none* as appropriate.
- [tpctool](#)
tpctool is a stand-alone Java™ client and connects to the Device server only. *Tpctool* connects through TCP, HTTP, and SOAP to the web service APIs. The commands provide query, control, and reporting capabilities only. The commands do not initiate discovery, probes, or configuration and control of agents.
- [Fabric-specific problems](#)
You can use this information to troubleshoot fabric-specific problems.
- [SMI-S providers](#)
The following information provides details on how to troubleshoot CIM client or SMI-S provider problems.
- [Linux SRA probes hanging problem](#)
The Storage Resource agent (SRA) probes use Logical Volume Manager (LVM) commands on Linux® that can hang for various reasons, for example, if a disk is suspended.

Discovery

Use this information to troubleshoot discovery problems.

General information for problem diagnosis

For discovery problems, there are several service level logs to check:

- msgTPCDeviceServer.log
- traceTPCDeviceServer.log
- dmSvcTrace.log

These logs are in this directory: *installation_dir*\IBM\TPC\device\log.

You can get a core dump file in this directory: *installation_dir*\ewas\profiles\DeviceServerProfile\.

For Windows operating systems, follow these steps:

1. Go to this directory: *installation_dir*\device\apps\was\bin
2. Enter the following command:

```
wsadmin set jvm [$AdminControl completeObjectName
type=JVM,process=deviceServer1,*] $AdminControl invoke $jvm dumpThreads
```

For Linux® or AIX® operating systems, enter this command:

```
ps -ef | grep "IBM/TPC" | grep "ewas" | grep "deviceServer" |
awk '{print $2}' | xargs kill -3
```

The **setenv** command must be run before you run the **srmcp** command.

For Windows operating systems, follow these steps:

1. Go to this directory: *installation_dir*\device\bin\w32-ix86
2. Enter the following command: **setenv**.

For Linux or AIX operating systems, follow these steps:

1. Go to this default directory:

Linux systems:
installation_dir/device/bin/linux

AIX systems:
installation_dir/device/bin/aix

2. Run this command:

```
./setenv.sh
```

To discover the job activity, issue the following command:

```
srmcp -u user_ID -p password DiscoverService list jobs
```

To determine the Db2® activity, issue the following Db2 commands:

```
db2 list applications show detail > dbListApplicationDetails.out
db2 update monitor switches using statement on lock on table
db2 connect to tpcdb
db2 get snapshot for all on tpcdb > dbsnap.out
```

Common user errors

A discovery or probe operation does not complete. This condition can be caused by the following situations:

- There is a lock contention when IBM Spectrum® Control accesses the database tables.
- DB2® does not return information from a query.
- The SMI-S provider starts returning information and then stops.

To work around this problem, stop and start the Data server and Device server. Also, use the service tool to collect log information. For information about the service tool, see [Service tool overview](#).

Unable to discover filers in a Windows environment

If you are unable to discover filers in a Windows environment, check the following items:

- Check the discovery log file:
 - For initial discovery, check the probe_000001.log file.
 - For discovery jobs, check the associated discovery log file.
- If the log file contains the following information, check the authentication information that is entered in the Filer Login panel:

```
DIS0001I FULL command selected.
STA0246I Discovery started
STA0339I Discovery started for domain xxx
NAA0041E Local login failed -- SRM\db2admin
NAA0041E Local login failed -- SRM\db2admin
STA0249I Sending results to server
STA0250I Server saved results
```

This information indicates:

- Authentication that is entered for a single filer is validated.
- The default login and password, which is used for multiple filers, are not validated.
- Verify that the SNMP community specified in the discovery job is the SNMP community that is used by the filer.
- Check the TPCUser.Discovery*.trace log file. If the trace log file does not exist, set the agent trace to the maximum value and rerun discovery. The trace file for discovery contains helpful information about hosts that are found on the domain, which host is contacted that uses SNMP, the SNMP community name, and Windows API calls errors.

Unable to discover filers in a Linux or AIX environment

Check the following items:

- Ensure that the agent can communicate with the filer and can mount a file system from the filer.
- Verify that the SNMP community name specified in the discovery job is the SNMP community that is used by the filer.
- Ensure that the mounts on the agent system are current and not stale.
- Check the TPCUser.Discovery*.trace log file. If the trace log file does not exist, set the agent trace to the maximum value and rerun discovery. The trace file for discovery contains helpful information about what SNMP community name is used to communicate with the filer.

Other scenarios that can cause discovery problems

Other scenarios that might occur are as follows:

- If you see message STA0155E in the Windows discovery log file:

```
DIS0001I FULL command selected.
STA0246I Discovery started
STA0339I Discovery started for domain xxx
STA0155E Discovery was unable to find the Windows computers in the domain
or workgroup.
STA0249I Sending results to server
STA0250I Server saved results
STA0252I Discovery completed successfully.
```

If you receive these messages, turn on tracing to generate a discovery trace file. Check the trace log for error code 6118. If you see code 6118, start the Computer Browser service on the Windows system, if the service is not turned on. Then submit discovery again.

- For hosts that are on different Windows subnets but have the same workgroup name, IBM Spectrum Control considers the workgroup is on the same subnet and excludes one subnet when discovery is run. To work around this issue, disable the Storage Resource agent on one subnet and then run discovery again.

Discovery and probe completion codes

Use this information to interpret status and completion codes returned in messages from discovery and probe jobs.

Status and return codes

Discovery or probe process completion status and return codes are referenced in messages, for example:

- BTADS0001I Discover Process with Device Server RUN ID 153001 and Job ID 405024 is complete with Status= 1 , Return Code= 0

- BTADS0102E The discover process that has the Device server run ID 112003 and job ID 291022 completed with one or more errors. The process status is 4 and the return code is 8

Table 1. Status codes

Code	Value
0	Failed
1	Successful
2	Running
3	Warning
4	Partial failure
5	Cancelled
-1	Submitted
-2	Waiting

Table 2. Return codes associated with completion codes

Code	Value
0	Success
4	Warning
8	Failed

Monitoring service

The following information shows you how to troubleshoot monitoring service problems.

IBM Spectrum® Control shows an event for the discovery and control components for the following actions:

- When you add a resource.
- When you make resource status changes.
- When the resource is not detected.

IBM Spectrum Control provides an alert generation on behalf of the performance monitoring service when the threshold changes and there is a collection failure. The database table snapshot for a certain resource type is based on the request that is made to the performance monitoring service.

For general diagnosis of the problem, check these items:

- Make sure that the alert trigger condition is set up correctly.
- Make sure that the changes did happen.
- Check the alerts that are shown on the Alerts page in the IBM Spectrum Control GUI.

For specific diagnosis of the problem, check these items:

- Locate the Device server trace log and search for possible exceptions that are generated from the monitor service.
- If no errors and exceptions are found for the monitor service, it means that the monitor service is running without problems. Next, check to see whether some other component is causing the problem such as the Change Detection component.

Common usability problems

These items are some common user errors with the monitoring service:

Why do I not get the correct alert?

Check to see that you configured the correct trigger condition for the correct resource or resource type.

Why is overall performance slow?

Many database activities to form SNMP and Netcool®/OMNIBus events might affect overall performance if too many alerts get created. You can configure only the needed alerts to control the number of alerts.

Why am I losing an event or alert?

An event or alert might be lost because the required data is unavailable in the database or there are network communication problems. After the alert creation fails, the alert for this changed device is not re-created.

Why do I see unexpected alert definitions and unexpected alerts with the keyword Pre-defined in the name?

These alerts are shown when another application is subscribed to receive events from IBM Spectrum Control through the product API.

You can change the setting to hide Pre-defined alert definitions when you use the **tpctool** command:

```
tpctool setdsconf -user user_ID -pwd password -url localhost:9550
-property APIAlert.Visibility false
```

Logs and traces

When a problem occurs, you can set the trace level parameter `san.eventFactoryTrace.level = INFO` when you use the CLI command to get detailed trace information. Make sure that the created FabricAlert object is sent to the Data Server successfully.

The detailed trace information is saved in the following directory: `install_directory\data\log\traceTPCDeviceServer.log`

Use the service tool and repocopy tool to collect trace information for IBM Spectrum Control and the database.

When an expected alert is not presented, check the following items:

- The trace log to make sure that the resource changes are reported to the monitor service. For example, search the log for the corresponding API call.

- The correct alert population routine is used, and any other exception that is encountered during the process.
- Make sure that the alert object is sent to the Device server.

Performance monitoring

The following information shows you how to troubleshoot performance monitoring problems.

Tuning performance monitors for switches

When performance monitors for switches encounter timeout problems, you can change the attributes that affect the performance monitor. The performance monitor uses algorithms to collect performance information. The association algorithm is optimized for environments with larger numbers of switches that are managed by an SMI Agent or SNMP agent. The enumeration algorithm is optimized for environments with fewer switches that are managed by an SMI Agent or SNMP agent. IBM Spectrum® Control attempts to determine which algorithm to use based on the fabric configuration.

You can change the following attributes in the pm.conf file:

MinSwitchPortRatio

This attribute defines the minimum port ratio, for example, the ratio of monitored ports to total ports for a switch data source. This value is used to determine which algorithm to use for querying the performance statistics data for the ports through the SMI Agent or SNMP agent. The default cutoff value for the enumeration algorithm is 20%. If the cutoff value is less than 20% of the ports for an agent that are targeted for performance data collection, the association algorithm is used instead, if the MaxSwitchPortLimit is not exceeded. This parameter does not apply to Cisco SNMP agents.

MaxSwitchPortLimit

This attribute defines the maximum port limit, for example, the maximum number of ports to be monitored that use the association algorithm for switch SMI Agents. This value is used to determine which algorithm to use for querying the performance statistics data for the ports through the SMI agent. If not set, the default cutoff value for the association algorithm is 256 ports, so that performance data is not collected for more than 256 ports simultaneously that use the association algorithm. Instead, the enumeration algorithm will be used for any SMI agents after the limit is reached.

To change these attributes in the pm.conf file, go to the following directory: *SC_installion_directory/device/conf/pm.conf* and remove the number sign (#) from these attributes and modify your setting:

```
#com.ibm.tpc.perf.MinSwitchPortRatio = 0.4  
#com.ibm.tpc.perf.MaxSwitchPortLimit = 256
```

Save the file. You must restart the Device server for these changes to take effect.

Configuration History

Diagnose problems with the Configuration History view.

The Configuration History view is a variation of the topology viewer. The Configuration History view shows the current configuration and provides information about what the configuration looked like in the past. Use the snapshot selection panel in the Configuration History page to view changes that occurred in your storage environment between two or more periods.

The Configuration History views can show a resource two or more times. This condition might happen when a resource is removed manually from the database or is removed automatically according to the Data for missing resources setting on the History Retention page, and IBM Spectrum® Control rediscovers that resource again. IBM Spectrum Control assigns the resource to a different ID when it is rediscovered.

Limitations

The following items are not supported by Configuration History:

- Alert and performance overlays
- Pin and refresh settings actions

Diagnosing problems

Check the following conditions when you encounter a problem with Configuration History:

- Message log for the Device server. The log file is available in the following directory: *installation_dir/device/log/msgTPCDeviceServer.log*. The message prefix is HWNCC.
- Device Server trace log.

To turn on tracing for Configuration History in the Device server, complete these steps:

1. In the menu bar, go to Home > System Management.
2. Click Component Servers in the Components section. In the Component Servers pane, you can view performance information for each component server. The current trace level for each component server and the Device server services is highlighted.
3. Adjust the trace level for a component server or a Device server service by clicking Low, Medium, or High.

The output data is logged in the traceTPCDeviceServer.log file.

- Data Server trace log. For information about turning on tracing for the Data server, see [Tracing servers and agents](#). Select High to collect the most data possible for this server. The Data server log files are at this location: *installation_dir/data/log/ServerTrace.log* and *installation_dir/data/log/dbtrace.log*.

SMI-S fabric probe

The following information shows you how to diagnose problems with the SMI-S probes of fabrics.

IBM Spectrum® Control supports the SMI-S agent (also called the CIM agent or CIMOM) for the fabric. The SMI-S agent collects basic fabric, switch, and port information that is used for performance monitoring. For Brocade fabrics, full topology information and zoning information is also collected during the fabric probe.

For a list of switches and directors that are supported by IBM Spectrum Control, see the [IBM Spectrum Control interoperability matrix for switches](#). Search for the appropriate version of switch and director.

Common usability problems

SMI-S fabric usability issues include the following items:

Switch Vendor and Model are blank or unknown

Some switch vendors do not provide a distinguishable vendor name or model name through some of the fabric agents.

Switch performance monitoring job fails immediately

Check to see that a probe of the related fabric completed.

The Switch port, Index, Slot, or Enabled State option is blank for the Brocade switch

Check to see that an SMI-S agent is configured for this fabric, and check to see that a fabric probe completed.

CIMOM discovery failed

Check to see that the switch SMI-S agents are at the supported levels.

Network Advisor SMI Agent

For information about the Network Advisor SMI agent, see the following documents:

- *Brocade Network Advisor User Manual* that supports your version of Network Advisor.
- *Brocade SMI Agent User's Guide* for the version of the SMI Agent you are using.
- *Brocade Network Advisor (BNA) Integrated SMI Agent* for any problems that are encountered with this type of SMI Agent.

Things to note

When you have a problem with the SMI-S fabric probe, check these items:

- The following information is collected only through the SMI-S agent:
 - Brocade blades
 - Brocade Switch FCPort EnabledState and Index
- When a discovery is performed, IBM Spectrum Control finds fabrics and switches through the switch SMI-S agents. IBM Spectrum Control finds basic information for switches that are not registered with the SMI-S agent, if they are in a fabric with a switch that is registered with the agent. IBM Spectrum Control also finds and persists fabrics and the fabric-to-switch relationships.
- When a fabric is probed, the following actions occur:
 - SMI-S agent discovery is rerun.
 - When a backend fabric agent assignment is made, IBM Spectrum Control checks to see which SMI-S agents are available. The agents are sorted and the top agent in the list is run. The other agents in the list are run only if the fabric probe fails with the previous agent.
- Switch blade port index values
 - The switch blade port index values might be incorrect for Brocade switches where one of the following scenarios apply:
 - When a port swap operation is performed.
 - When the Extended Edge PID mode has been enabled on the director. For more information about the Extended Edge PID mode, see the Brocade document *Fabric OS Administrator's Guide*.

VMware ESX

The following information shows you how to diagnose problems with VMware ESX servers.

IBM Spectrum® Control supports the VMware Virtual Infrastructure that consists of the ESX server and VMware VirtualCenter. The ESX server is a true hypervisor product that can host multiple virtual machines that run independently of each other while they share hardware resources.

IBM Spectrum Control supports the following services:

- The Virtual Center discovery is the only successful criteria.
- Alerts are for probe only. No event, traps, or alarms from the Virtual Infrastructure are supported.
- IBM Spectrum Control supports reporting only. No active management of the Virtual Infrastructure is supported.
- The information that is obtained from VMware is equivalent to what is obtained through a read-only Storage Resource agent. It does not include scripts or file system (FS) extension.

To get detailed information for space and capacity reports, you must complete the following steps:

- Probe all the ESX servers.
- Install Storage Resource agents on all the virtual machines and probe them. Otherwise, you get a partial picture with incomplete capacity data.

Common usability problems

The following list shows some common user errors with VMware:

- There are new totals in the dashboard and reports.
- Note the indications for the hypervisor status.
- Differentiate Storage Resource agent for a virtual machine information from the ESX Server.

Dependencies

Some of the dependencies are as follows:

- VMware ESX 3.0.1 and Virtual Center 2.0.1 are required.
- A probe of the hypervisor is required for complete reporting.
- Hypervisors and Virtual Centers have HTTPS communication that is turned on and HTTP communication that is turned off by default.
- You require permission to browse through the data stores on the ESX server. For more information about how to check permissions to browse data stores, search the IBM Spectrum Control Knowledge documentation for *hypervisor permissions*.

Diagnosing the problem

When you troubleshoot a VMware problem, make a note of the following items: When you troubleshoot problems with VMware devices, make a note of the following items:

- View the job logs for probe create jobs that might indicate that tracing information is available.
- Look for trace messages in the general trace file: *installation_dir/device/log/traceTPCDeviceServer.log*
- Enable tracing for the Device server.
 1. In the menu bar, go to Home > System Management.
 2. Click Component Servers in the Components section.
 3. In the Device section, click On for each of the services of the Device server.

The output data is logged in the traceTPCDeviceServer.log file.

For more information about setting the trace levels of the component servers, see the online help for the System Management page.

FlashCopy

IBM Spectrum® Control FlashCopy® support labels volumes as having a FlashCopy property of *source*, *target*, or *none* as appropriate.

IBM Spectrum Control does not include the ability to create or manipulate FlashCopy relationships.

Diagnosing the problem

Most issues arise from one of two causes:

Missing data

The storage system does not report the FlashCopy data in a manner that IBM Spectrum Control expects. Missing data is most likely to occur in the event of the release of new microcode or CIM agents for a storage system.

For storage systems that use CIM agents, use a CIMOM browser to verify that those agents are aware of the FlashCopy relationships.

Tip: IBM® storage systems, such as DS8000® and XIV®, do not require CIM agents. These storage systems use native interfaces.

Improperly created FlashCopy relationships

A persistent FlashCopy relationship was not created. If the relationship is temporary, the relationship is probably no longer valid by the time IBM Spectrum Control probe requests the information. Verify that the management software of the storage system reports the FlashCopy relationship.

tpctool

tpctool is a stand-alone Java™ client and connects to the Device server only. *Tpctool* connects through TCP, HTTP, and SOAP to the web service APIs. The commands provide query, control, and reporting capabilities only. The commands do not initiate discovery, probes, or configuration and control of agents.

Tpctool is installed in the following default directories:

Windows operating systems:

installation_dir\cli

Linux® or AIX® operating systems:

installation_dir/cli

The Windows command to run the tool is *tpctool*. The Linux or AIX command to run the tool is **tpctool.sh**.

There are two distinct kinds of authentication:

- User authentication
- Super user or host-based authentication

The user authentication requires a user ID and password authenticated in the Device server authentication domain. Role-based authorization is enforced on a per command basis.

The super user or host authentication is for the tpc_superuser user ID. The password for this user ID bypasses role-based authorization. This action is the required authentication method for AIX-based Device servers.

Some of the control commands run for a long time. An example is the **mkvo1** command. It is not possible to determine the intermediate status of the command. If the *tpctool* client stops, for example the user presses Ctrl-C or the node crashes, all connections with the job are lost. It is not possible to reconnect to the host to check on the status of the command.

The Device server logs and trace files, assuming that tracing is on, are shown. For the disk commands, see this log:

DiskManagerService: dmSvcTrace.log

For the Fabric commands, see these logs for the Fabric Manager Service:

- TPCZoneControl.log
- msgTPCDeviceServer.log
- traceTPCDeviceServer.log

For the reporting commands, see this log:

PerformanceService: tracePerfMgr.log

For fabric reports, the commands pass through the Fabric service for authorization. For subsystem reports, the commands pass through the Disk service for authorization. For configuration commands, see this log:

ConfigService: traceTPCDeviceServer.log

For all other commands, see this log:

msgTPCDeviceServer.log

Here are some notes about passwords:

- You can use the GUI to change the Device server host authentication password.
- If you use the GUI to change the password, the password is updated in the database and the Device server is notified.
- The configuration file for all the Fabric agents must be manually changed.

Common usability problems

Here are some common usability problems with the *tpctool*:

- You do not know where the installation directory is. Is the installation directory in a nonstandard location?
- You did not set the *PATH* or *chdir* to the *installation_dir/cli* directory.
- Using *tpctool* on Linux or AIX operating systems instead of *tpctool.sh*.
- Did not provide the following parameters in the command:
 - **-url** *tpcserver:deviceserverport* where *tpcserver* is the hostname or IP address of the server that runs IBM Spectrum® Control and the *deviceserverport* default value is 9550, for example: localhost:9550
 - **-user**
 - **-pwd**

A common error message for *tpctool* is as follows:

AAJ000009E Error communicating to the App server.

This error message indicates an invalid port, invalid host, or that the Device server is unreachable. To distinguish what the problem is, you can set a debug flag that prints the Java stack if an exception occurs. Most errors are propagated to *tpctool* as exceptions. To set this flag:

Windows operating systems:

set TPCCLIDBG=1

Linux or AIX operating systems:

export TPCCLIDBG=1

For an invalid host, the host name is embedded in the message. For example:

```
[SOAPException: faultCode=SOAP-ENV:Client: msg=Error opening socket:
java.net.UnknownHostException:badhost:...]
```

For a valid host but an invalid port or when the Device server is down, an example of a message is as follows:

```
[SOAPException: faultCode=SOAP-ENV:Client: msg=Error opening socket:
java.net.ConnectException: Connection refused:...]
```

An invalid port is the most common problem.

If the Device server node is unreachable, an example of the message is as follows:

```
[SOAPException: faultCode=SOAP-ENV:Client: msg=Error opening socket:
java.net.SocketException: Operation timed out:
connect:could be due to invalid address:...]
```

srmcp commands

Before you run a **srmcp** command, you must run the **setenv** command.

Windows operating systems:

Follow these steps:

1. Go to this default directory: *installation_dir\device\bin\w32-ix86*
2. Run the following command: **setenv**.

Linux or AIX operating systems:

Follow these steps:

1. Go to this default directory on UNIX

Linux systems:

installation_dir/device/bin/linux

AIX systems:

installation_dir/device/bin/aix

2. Run this command:

```
./setenv.sh
```

You can use the following **srncmp** commands for SANEventCorrelatorFactory for fabric troubleshooting:

- To list the configured filters:

```
cd installation_dir\device\bin\w32-ix86
srncmp -u user_ID -p password SANEventCorrelatorFactory
list
```

- To add a filter with the specified values (any or all can be specified). You can specify just the enterprise, the OIDs, or a combination:

```
cd installation_dir\device\bin\w32-ix86
srncmp -u user_ID -p password SANEventCorrelatorFactory
add filter [enterprise=enterprise] [genericTrap=<number>]
[specificTrap=number] [OID=value]
```

- To specify an IP address that is filtered for the specified filter:

```
cd installation_dir\device\bin\w32-ix86
srncmp -u user_ID -p password SANEventCorrelatorFactory
add address ID IP_address
```

- To remove a specified filter:

```
cd installation_dir\device\bin\w32-ix86
srncmp -u user_ID -p password SANEventCorrelatorFactory
remove filter ID
```

- To remove the IP address from the filter list:

```
cd installation_dir\device\bin\w32-ix86
srncmp -u user_ID -p password SANEventCorrelatorFactory
remove address ID IP_address
```

Here is the **srncmp** command to change the host authentication password:

```
cd installation_dir\device\bin\w32-ix86
srncmp -u user_ID -p password ConfigService setAuthenticationPw
new_host_password
```

Here is the **srncmp** command to change the DB2® password (not the db2admin password) that the server uses:

```
cd installation_dir\device\bin\w32-ix86
srncmp -u user_ID -p password ConfigService setPw
new_host_password
```

Fabric-specific problems

You can use this information to troubleshoot fabric-specific problems.

Fabric discovery

You might encounter the following of problems when you discover the fabrics in a storage environment:

- An error in the logs. Typically the first error you see is the most important.
- Information in the GUI that is incorrect or unexpected.

For information about error messages, see [Message types](#). Follow the explanations, user responses, and administrator responses.

For problems related to the discovery of fabrics, follow these steps:

- Run the service tool. See [Service tool overview](#).
- Run the **repocopy** command to capture the IBM Spectrum Control database. Follow these steps:
 - Go to the following directory:

Windows operating systems:
`installation_directory\data\server\tools`

Linux® or AIX® operating systems:
`installation_directory/data/server/tools`

- Run the following command:

Windows operating systems:
`repocopy.bat`

Linux or AIX operating systems:
`repocopy.sh`

- Select Export data from repository tables.
- Select the directory for the export file location.
- Use the default values for delimiter and double quotation marks.
- Use the default values for the connection properties.

- g. Compress all the files in the output directory and send to the IBM Software Support.
3. Get more trace information. To turn on or change the tracing for the component servers and Device server services, follow these steps:
 - a. Go to Home > System Management.
 - b. Click Component Servers in the Components section. In the Component Servers pane, you can view performance information for each component server. The current trace level for each component server and the Device server services is highlighted.
 - c. Adjust the trace level for a component server or a Device server service by clicking Low, Medium, or High.
4. Correlate the time stamps in the job log.
5. Check the health of the Device server.
6. Check these files:
 - `installation_directory\device\log\traceTPCDeviceServer.log`
 - `installation_directory\device\log\msgTPCDeviceServer.log`

To check on incorrect information, check these items:

- Did IBM Spectrum Control receive an event? Look for the following items:
 - Events in the alerts.
 - ALR4100I: Received an SNMP trap notification from source.
- The health of the probe:
 - Look in the log file: `installation_directory\device\log\msg.probeFabricAgents.x.x.log`
 - Look for exceptions or error messages.
- If the population of the results into DB2® failed:
 - Check the following log file: `installation_directory\device\log\msg.probeFabricAgents.x.x.log`
 - Look for exceptions or error messages.

The following is a list of some of the known problems:

- Some unsupported devices like multi-protocol routers can cause failures during the discovery process.
- AIX HBA information is limited and provides only name, manufacturer, firmware, and model.
- HBA hardware and BIOS versions are not shown in the IBM Spectrum Control GUI.

HBAs not discovered by a Storage Resource agent

The Storage Resource agents do not use RNID and SCSI inquiries for remote HBA and resource identification. You need a Storage Resource agent on each system for which HBA information is needed. You need a native interface storage system or SIM-S Provider to identify the storage devices.

Check for any messages in the probe log for information about what might be the problem. You can also check the following items:

- If the supported HBA driver and API are installed. In some cases, the HBA driver includes the API. In other cases, the HBA driver and API must be installed separately. Contact your HBA vendor for information about the driver and API.
- If the HBA and API are at a level that is supported by IBM Spectrum Control. See <https://www.ibm.com/support/pages/node/388393>.
- There are some operating system restrictions for Storage Resource agent functions.

For more information about the requirements for the Storage Resource agent, see <https://www.ibm.com/support/pages/node/6249361#Agents>.

- Check to see whether the HBA is configured appropriately. For information about how to configure the HBA, contact your HBA vendor.

The HBA to port relationship is a logical concept. Sometimes one physical card might have more than one HBA port. Multiple ports in the same HBA share node WWN.

Probe failures

If the probe of a fabric fails, check the following items:

- Look at the probe job log and Device server message and trace log files. If the error message is in processing before the agent is called, the issue is with the agent assignment.
- Check for communication errors when the agent is called, any error response from the agent, or no response from the agent.
- Check the agent log files for error messages.

Fabric configuration

Check the following items for fabric configuration problems:

- Look at the GUI or CLI error code. See IBM Spectrum Control Messages for the error codes. The error codes have two parts: command failure code and native failure code.

Command failure code

The command failure code specifies which operation failed, such as an error that occurred while a zone is created.

Native failure code

The native failure code gives the reason for the error, such as *Already exists*.

- Whether the agent that is connected to the fabric is configured appropriately for zoning.

Check the following logs:

- Web server log files.
- `AuditTrace.log` in the Data Server log directory.
- `TPCZoneControl.log` in the Device server log directory.
- `msg.control.123.123.log` in the Device server log directory
- `traceTPCDeviceServer.log` in the Device server log directory

How to adjust log and trace levels

You can adjust the following trace levels in the `DataStore.properties` and `native.log.properties` files in the Device server conf directory.
Tip: Make a backup copy of the `DataStore.properties` file before you modify this file. A corrupt `DataStore.properties` file can cause IBM Spectrum Control to not start.

```
san.ControlTraceLogger.level=ALL
san.SanZoneControlTraceLogger.level=ALL
san.SanZoneControlAgentTraceLogger.level=ALL
```

Commands and tools for troubleshooting

- On the seed switch of the fabric, check for a transaction lock. For example, for the Brocade CLI, use the `cfgTransShow` and `cfgTransAbort` commands.
- Look for IBM Spectrum Control GUI or CLI instances that can be holding a transaction lock. It might be difficult to abandon the transaction lock on older switch models without rebooting the switch.

Fabric events

Some of the common reasons for not getting notification of external fabric events or not having fabric probes run automatically in response to events are as follows:

- Switches are not configured to send SNMP traps to the server.
- The fabric for which an event was received or fabrics that are discovered by the Storage Resource agent are not included in any probes that are defined by the user. An automatic probe of fabrics is done only for fabrics that are included in some probe definition.
- The switch is configured for a trap level that results in the filtering of the trap and not sent from the switch to the server. The trap that is sent from the switch to the server is filtered by the server because of default filter rules or filtering rules, or both, configured by the user.

To configure tracing to check that the traps are filtered, set the trace level parameter `dsn.eventFactoryTrace.level = INFO` by using the CLI command to get detailed trace information. Ensure that the created FabricAlert object is sent to the Data Server successfully.

The detailed trace information is saved in the following directory: `installation_directory\data\log\traceTPCDeviceServer.log`.

If you are not getting a fabric alert when you expected one, check if the information shown in IBM Spectrum Control reflects the change. Check that the appropriate types of alerts are defined for the fabric.

Check the following logs for error messages:

- `installation_directory\data\log\server_xxxx.log`
- `installation_directory\data\log\TPCD_xxxx.log`

CIM indications

Here is a list of some things to check when you are having issues with CIM indications from Network Advisor:

CIM indication not received

- Verify that SNMP registration is successful by the Network Advisor discovery dialog
- Verify if the configured SNMP port is blocked by a firewall
- Verify if the Management Information Base (MIB) capability for trap forwarding and tracking changes are configured properly

To get the Network Advisor logs, follow these steps:

1. Open the Network Advisor Server Management Console.
2. Go to the Technical Support Information tab.
3. Click Browse and select the log destination. Click capture.

Fabric data retention

Check for problems that are related to the removal of missing fabric resources and data in the following logs:

- `server_xxxxx.log`
- `TPCD_xxxxx.log`

These logs are in the following directory: `installation_directory\data\log\`

Related reference

- [Default locations of log files](#)

SMI-S providers

The following information provides details on how to troubleshoot CIM client or SMI-S provider problems.

XML requests and responses are written to the CIM client trace file that is in the following directory: `C:\Program Files\IBM\TPC\device\log\cim-xml-trace.txt`

The `cimxmlTracing` configuration parameter in the Db2® table must be set to true to collect tracing information.

To set the CIM client trace parameter `cimxmlTracing`, follow these steps:

1. Check the value of the `cimxmlTracing` attribute in the Db2 table by using the following command:

```
tpctool getdscfg -user user_ID -pwd password -url localhost:9550
-property cimxmlTracing
```

The `tpctool` command shows the following information:

Property	Context	Value
cimxmlTracing	CIM	false

If the value is false, set the attribute to true:

```
tpctool setdscfg -user user_ID -pwd password -url localhost:9550  
-context CIM -property cimxmlTracing true
```

Run the `getdscfg` command again to show the new value of the attribute:

Property	Context	Value
cimxmlTracing	CIM	true

2. After you modify the value of the attribute, stop and restart the Device server.

Linux SRA probes hanging problem

The Storage Resource agent (SRA) probes use Logical Volume Manager (LVM) commands on Linux® that can hang for various reasons, for example, if a disk is suspended.

Problem

When a Storage Resource agent is deployed on a Linux host, it starts several Linux commands, including **lvm pvs**. The Storage Resource agent probe can hang at the same point that the **lvm** command stopped working.

On the IBM Spectrum® Control web user interface, the Storage Resource agent probe remains in an In Progress state and does not give any indication that the probe is hung.

Action

To verify that the **lvm** command is causing the SRA to hang, look at the SRA trace log on the SRA deployment and check that the last command run is an **lvm** command.

To solve the issue, follow these steps:

1. Kill the hung SRA process.
2. Kill the hung **lvm** command.
3. Rerun the **lvm pvs** command from the command line.
4. Verify that the **lvm** command runs successfully by running the **pvdisk** command to display the volume details.
5. Restart the SRA deployment.

Related information

- [fqz0 t enabling a data agent.html#fqz0 t enabling a data agent](#)

Getting support

For help with resolving issues with IBM Spectrum® Control, you can contact IBM® Support or use IBM self-help resources.

To get help from IBM Support and report issues with IBM Spectrum Control, complete these steps:

1. Log in to the [IBM Support portal](#). If you don't have an IBM ID, go to [Create your IBM account](#) and complete the form.
2. Click Open a case.
3. Complete the form.
4. Click Submit case.

Optionally, you can contact IBM Support by phone. In the U.S., call 1-800-426-4968. For other countries, go to <https://www.ibm.com/planetwide/>. When you open a case or contact IBM Support by phone, have the following information ready:

- The version, release, modification, and service level number of IBM Spectrum Control that you're using.
- The communication protocol (for example, TCP/IP), version, and release number that you're using.
- The activity that you were doing when the problem occurred, including the steps that you followed before the problem occurred.
- The exact text of any error messages.

Videos about the IBM Support portal: Watch a few short videos to learn more about the IBM Support portal:

- [Introducing a new customer portal](#)
- [Open and manage cases](#)
- [IBM Support Community: Search](#)
- [IBM Support Community: Forums](#)

Collecting and uploading service logs

To help close your case, IBM Support might need you to collect and upload your service logs. To complete this task, use one of the following methods:

- Package and send the logs from the System Management page. See [Packaging and sending log files from the System Management page](#).
- Package the logs from the command line and then manually upload the resulting .zip file. See [Packaging log files from the command line and sending them to IBM Support](#).

Using IBM self-help resources

Use the following resources to learn about and help troubleshoot issues that you might encounter in IBM Spectrum Control.

[IBM Support page for IBM Spectrum Control](#)

View, start, or contribute to IBM Spectrum Control community discussions, view supplemental resources, or chat with IBM Watson®.

[Fixes](#)

Find and download the latest fixes and updates for IBM Spectrum Control.

[Product updates and security fixes](#)

Find information about IBM Spectrum Control releases and downloads, security, subscribing to IBM announcements, collaborating with the IBM Spectrum Control team, reporting problems, and End of Support (EOS) details for specific releases of IBM Spectrum Control and related products.

[Social media](#)

Watch videos and read blogs to learn more about how to use IBM Spectrum Control to manage your storage environment.

Messages

View details about error, warning, and informational messages in IBM Spectrum® Control.

- [Introduction to messages](#)
The message ID can tell you more about the type of information you are dealing with.
- [Message types](#)
Read detailed information on specific messages.

Introduction to messages

The message ID can tell you more about the type of information you are dealing with.

Messages may be displayed onscreen or logged in text files. The log files for the Data Server are located in the `\install_dir\data\log` directory and the log files for the Device server are located in the `\install_dir\device\log` directory on the manager machine. You can optionally configure the maximum retention period for each log file.

The following example illustrates the message ID format:

BPCUI0053E

Where:

BPCUI

A 3- to 5-character prefix that identifies the message group.

0053

A 4- to 6-number message identifier.

E

The message type:

E

Error. An Error message indicates a situation where you need to take remedial action. The message help provides an explanation and suggests possible action.

W

Warning. A Warning message alerts you to a situation that may need your attention. The message help provides an explanation and suggests action or pointers for investigation.

I

Information. An Information message provides additional detail on the outcome of an action. No further action is required.

Message text variables are displayed in *italics*.

Message groups

The alphabetic identifier at the start of message IDs indicates their function.

Identifier	Function
AGT	Agent
ALR	Alert
BPCCA	Collector agent
BPCCM	Collector Managers
BPCDP	Data Processor
BPCIN	On premise installation
BPCRE	Alert server
BPCRS	Middleware
BPCSS	Scheduler
BPCUI	Web client
BTACD	Database verifier, SAN database service

Identifier	Function
BTACE	SAN event service, messaging middleware
BTACS	Service manager
BTADS	Fabric discovery
BTAEC	Event correlator factory
BTAHM	Host manager
BTAIC	In-band change agent
BTALG	Logging service
BTAMS	Messaging service
BTAQE	Query engine
BTASA	Scanner agent
BTASD	SAN manager daemon, fabric user interface
BTATG	UNIX CLI help
BTAVM	Hypervisor management
BTAZC	Zone control
BTM	Common Information Model Agent
BWN	Disk user interface
CMMNP	CLI infrastructure
CMMUI	CLI infrastructure
CNFG	Configuration
DIS	Discovery
EMSG	DS8000 master console
GEN	General
GPC	Performance User Interface
HWNAS	Agentless server
HWNAL	Authentication
HWNCA	Social, Mobile, Analytics and Cloud (SMAC)
HWNDA	Storage Resource Group
HWNEM	Element manager management
HWNFS	File system monitor
HWNLM	Host planner, security planner, and subsystem planner
HWNOP	Storage Optimizer
HWNPM	Performance Manager
HWNRM	Replication Manager
HWNSS	Single sign-on
JSS	Database
NAD	Storage Resource agents
NAG, NAS	NAS
SAA, SAG, SAS	SAN (some SAS messages are related to file system extension)
SRV	Server
STA, STG, STS	Base product (some STA and STS messages are related to file system extension)
VPLG	VASA provider

Message types

Read detailed information on specific messages.

- [AGT - Storage Resource Agent messages](#)
- [ALR - Spectrum Control Alert messages](#)
- [BPCCA - Data collector installation messages](#)
- [BPCIN - Spectrum Control installation messages](#)
- [BPCRE - Alert server messages](#)
- [BPCRS - Spectrum Control Middleware messages](#)
- [BPCCM - Data collector messages](#)
- [BPCDP - Data processor messages](#)
- [BPCSS - Scheduler messages](#)
- [BPCUI - User Interface messages](#)
- [BTACD - Database verifier messages for SAN database service](#)
- [BTACE - SAN event services messages](#)
- [BTACS - Service manager messages](#)
- [BTADS/BTAFM/BTAVM/HWN - Job logging messages](#)
- [BTAEC - Event correlator messages](#)
- [BTAHM - Host manager messages](#)
- [BTAIC - Inband change agent messages](#)
- [BTALG - Logging toolkit messages](#)
- [BTAMS - Spectrum Control Messaging Service messages](#)
- [BTAQE - Spectrum Control Query Engine messages](#)
- [BTASA - Spectrum Control SAN scanner agent messages](#)
- [BTASD - Fabric User Interface messages](#)
- [BTATG - UNIX Command Line Interface \(CLI\) help messages](#)

- [BTAVM](#)
- [BTAZC - Zone control agent messages](#)
- [BTM - Common Information Model \(CIM\) agent messages](#)
- [BWN - Disk User Interface messages](#)
- [CMMNP - Command Line Interface \(CLI\) infrastructure messages](#)
- [CMMUI - CIM Object Manager messages](#)
- [CNFG - Spectrum Control Configuration messages](#)
- [DIS - Discovery messages](#)
- [EMSG - DS8000 management console messages](#)
- [GEN - General Spectrum Control messages](#)
- [GPC - Performance User Interface messages](#)
- [HWNAS - Agentless Server messages](#)
- [HWNAB - Single sign-on service messages](#)
- [HWNCA - Storage multiple access control messages](#)
- [HWNDA - Data Manager API messages](#)
- [HWNEM - Element manager messages](#)
- [HWNFS - File system monitor messages](#)
- [HWNLM - Planner manager messages](#)
- [HWNOP - Storage optimizer messages](#)
- [HWNPM - Performance manager messages](#)
- [HWNRM - Replication manager messages](#)
- [HWNSS - Single sign-on User Interface messages](#)
- [JSS - Database messages](#)
- [NAD - Storage Resource Agent messages](#)
- [NAG - Storage Agent Resource messages](#)
- [NAS - Storage Agent Resource messages](#)
- [SAA - Storage Resource Agent - Storage Subsystem messages](#)
- [VPLG - VASA provider messages](#)

AGT - Storage Resource Agent messages

- [AGT0001E Cannot set working directory to directory.](#)
- [AGT0002E Invalid invocation: swtchusr.](#)
- [AGT0003I Agent registered.](#)
- [AGT0004I Agent started.](#)
- [AGT0005I Waiting for swtchusr.](#)
- [AGT0006E Process process number is already connected to the server.](#)
- [AGT0007E Another process is already connected to the server.](#)
- [AGT0008E Job-type should be job type, but it is job type.](#)
- [AGT0009E Error writing to swtchusr.](#)
- [AGT0010E Class not found: class name.](#)
- [AGT0011E class name does not inherit from class name.](#)
- [AGT0012E class name does not have appropriate constructor.](#)
- [AGT0013E Error constructing class name.](#)
- [AGT0014E Cannot close the log-file log file name.](#)
- [AGT0015E The process identified by the process number failed to start generating the error status number.](#)
- [AGT0016E Read from pipe failed.](#)
- [AGT0017E Write to pipe failed.](#)
- [AGT0018E Close of pipe failed.](#)
- [AGT0019W Problem reading GUID identified by the GUID number.](#)
- [AGT0020E Unparseable class-loader string class loader.](#)
- [AGT0021E Error stopping the bundle.](#)
- [AGT0031E User user name does not exist.](#)
- [AGT0032E Cannot exec user command shell shell -- errno = arguments error message.](#)
- [AGT0033E initgroups command\(shell\) fails -- errno = arguments error message.](#)
- [AGT0034E Cannot create log-file in directory -- errno = arguments error message.](#)
- [AGT0035W Error reading logical volume.](#)
- [AGT0036W Error deserializing from file name.](#)
- [AGT0037W Premature end of file -- file name.](#)
- [AGT0038W Class class name not found restoring from file name.](#)
- [AGT0039W Object restored from file name is not class name, Object is class name.](#)
- [AGT0040E Agent Shutting down.](#)
- [AGT0041E Agent not registered.](#)
- [AGT0042E Error writing file name.](#)
- [AGT0043E Error serializing to file name.](#)
- [AGT0044E Fatal error -- cannot connect to self port number.](#)
- [AGT0045E Error parsing configuration file on line line number.](#)
- [AGT0046W In agent.config, key name is not an integer -- default of default key number used.](#)
- [AGT0047E In agent.config, key name is not a valid port.](#)
- [AGT0048E In agent.config, no value supplied for key name.](#)
- [AGT0049E Too many unreported jobs.](#)
- [AGT0050E Error parsing configuration file.](#)
- [AGT0051W The configuration file had at least one bad value.](#)
- [AGT0052W Cannot create temporary file in directory.](#)
- [AGT0053W Tried to send signal process number to nonexistent PID process number.](#)
- [AGT0054W No running job \(job name, job number\).](#)
- [AGT0055E Jobs are not allowed to run as root.](#)

- [AGT0056I NOTICE: Server moved to computer name:port number.](#)
- [AGT0057E Unable to create message-queue from key value.](#)
- [AGT0058E Error waiting for job to start.](#)
- [AGT0059I Received request to shut down request type.](#)
- [AGT0060I Exiting normally.](#)
- [AGT0061I Exiting abnormally.](#)
- [AGT0062E Cannot fork job.](#)
- [AGT0063E Error terminating process.](#)
- [AGT0064E Putative SID has too few hyphens: security identifier.](#)
- [AGT0065E Error waiting for shutdown request with Queue ID.](#)
- [AGT0066E Error reaping.](#)
- [AGT0067E Error waiting for process process name.](#)
- [AGT0068E Cannot fork -- command: command name.](#)
- [AGT0069E open\(/proc/uptime\) fails.](#)
- [AGT0070E read\(/proc/uptime\) fails.](#)
- [AGT0071E /proc/uptime does not look right.](#)
- [AGT0072E Cannot determine system-boot time.](#)
- [AGT0073E Problem creating up-time poller.](#)
- [AGT0074E Unsupported Operating System.](#)
- [AGT0075E Cannot query key name key value.](#)
- [AGT0076E Performance title data title not found.](#)
- [AGT0077E <<System>> performance object not found.](#)
- [AGT0078E The <<system>> performance counter not found.](#)
- [AGT0079E Cannot create semaphore.](#)
- [AGT0080W Cannot retrieve process times.](#)
- [AGT0081W Cannot retrieve process exit code.](#)
- [AGT0082E Cannot log in user user name.](#)
- [AGT0083E Cannot obtain SID of local computer \(computer name\).](#)
- [AGT0084E Cannot get name of local computer.](#)
- [AGT0085E SID of local computer has unexpected type \(security identifier type\).](#)
- [AGT0086W The GetTokenInformation method failed to execute.](#)
- [AGT0087W The LookupAccountSid method failed to execute.](#)
- [AGT0088W Cannot find any domain controller for domain domain name.](#)
- [AGT0089W Cannot fetch the information for user domain\user.](#)
- [AGT0090W Cannot load profile for user username.](#)
- [AGT0091W Cannot unload user profile.](#)
- [AGT0092E Unable to retrieve status of job with the job number.](#)
- [AGT0093E Cannot retrieve environment block.](#)
- [AGT0094E Cannot impersonate to create resource.](#)
- [AGT0095W Environment for user username not fully configured.](#)
- [AGT0096E Malformed command-line.](#)
- [AGT0097E Command not found: fdisk command.](#)
- [AGT0098E Command has unsupported extension.](#)
- [AGT0099E Cannot duplicate handle.](#)
- [AGT0101E Cannot determine whether this is a domain controller.](#)
- [AGT0102E Cannot determine name of own domain.](#)
- [AGT0103W Cannot make Windows Job Object.](#)
- [AGT0104E Wrong server.](#)
- [AGT0105W Ignoring server relocation in config file.](#)
- [AGT0106I Server has changed.](#)
- [AGT0110E Cannot open key key name key value.](#)
- [AGT0111I Rereading config file.](#)
- [AGT0112E This product is not fully installed. To try again, stop and restart this agent.](#)
- [AGT0113E Cannot create temporary file in directory error message.](#)
- [AGT0114E Cannot write to file file name error message.](#)
- [AGT0115E Fork failed.](#)
- [AGT0116E Cannot exec command name.](#)
- [AGT0117W Cannot open file for auto-delete.](#)
- [AGT0118E Upgrader is not okay.](#)
- [AGT0119I version, modification and release.](#)
- [AGT0120E Error transmitting shutdown request to agent.](#)
- [AGT0121E Unable to create socket.](#)
- [AGT0122I The agent is already down.](#)
- [AGT0123W Unable to determine if agent is active. Sending shutdown request.](#)
- [AGT0124E Unable to create restartable job directory: directory name.](#)
- [AGT0125E Unable to read directory directory name.](#)
- [AGT0126E Unable to restart job from file file name.](#)
- [AGT0127E Unable to restart job job name.name run number run number.](#)
- [AGT0128W Cannot retrieve global structure.](#)
- [AGT0129W A system call failed in the agent program.](#)
- [AGT0130E The upgrader program parameters are missing.](#)
- [AGT0131I Exit Status = exit message.](#)
- [AGT0132E Cannot find/create script <script name>.](#)
- [AGT0133I Running Command: script name script contents.](#)
- [AGT0134E Putative SID does not start correctly: security identifier.](#)
- [AGT0135E Component component name is not a number: component number.](#)
- [AGT0136E No INSTANCEn.DAT file found!](#)
- [AGT0137E Cannot determine working directory!](#)
- [AGT0138W product name is not licensed on this computer.](#)
- [AGT0139E An initial probe cannot be performed.](#)

- [AGT0140W Discovery will not be performed.](#)
- [AGT0141W A filesystem scan will not be performed.](#)
- [AGT0142E Cannot make handle inheritable.](#)
- [AGT0143I file_count files scanned](#)
- [AGT0144I file_count total files scanned](#)
- [AGT0145I Retrieving job definition from server](#)
- [AGT0146I Scan started](#)
- [AGT0147I Retrieving report partition partition_number](#)
- [AGT0148I Report data retrieved](#)
- [AGT0149I Retrieving history data](#)
- [AGT0150I Deleting temporary files](#)
- [AGT0151E Unable to retrieve report definition](#)
- [AGT0152I Job definition retrieved](#)
- [AGT0153E MSCSEventListener thread interrupted. Shutting down native MSCS event listener.](#)
- [AGT0154E No IP address could be found for the local host.](#)
- [AGT0155I Windows Scan Option : scan parameters.](#)
- [AGT0156W Error killing process process_id.](#)
- [AGT0157E Interrupted\(agent shutting down\).](#)
- [AGT0159E A problem was encountered stopping IBM Spectrum Control subagent.](#)
- [AGT0160E Error enumerating keys under registry path.](#)
- [AGT0161W Error querying value root path\subkey_name\library file.](#)
- [AGT0162E Library : HBA_RegisterLibrary returned value.](#)
- [AGT0163E Library : HBA_LoadLibrary returned value.](#)
- [AGT0164W Library : HBA_GetAdapterName\(adapter number\) returned adapter name.](#)
- [AGT0165W Library : HBA_OpenAdapter\(adapter name\) failed.](#)
- [AGT0166W The HBA API HBA_GetAdapterAttributes for adapter adapter name returned error error code.](#)
- [AGT0167W Library : HBA_GetAdapterPortAttributes\(adapter number, port\) returned status.](#)
- [AGT0168W HBA data cannot be collected because the adapter name> adapter on the server does not support the HBA_GetAdapterAttributes function.](#)
- [AGT0200E Error waiting for process.](#)
- [AGT0250E Error starting bundle: exception message.](#)
- [AGT0251E Failed to install language pack.](#)
- [AGT0252E Failed to install language pack: File not found archive](#)
- [AGT0253E Invalid file format: file name](#)
- [AGT0254E Failed to load language pack.](#)
- [AGT0256I Waiting for Common Agent services.](#)
- [AGT0259E Agent cannot upgrade, the required space of 50 meg was not met.](#)
- [AGT0260E Failed to create a session with Reliable Scalable Cluster Technology Error ID-Error type: Error Description](#)
- [AGT0261E Failed to end a session with Reliable Scalable Cluster Technology Error ID-Error type: Error Description](#)
- [AGT0262W Concurrent cluster resource groups are not supported. The cluster will not be probed.](#)
- [AGT0263E Failed to get resource data from the HACMP cluster using the RSCT RMC interface.](#)
- [AGT0264E Failed to determine if the local node is clustered using the following command: command line](#)
- [AGT0265E Failed to get the cluster name using the following command: command line](#)
- [AGT0266E Failed to get the cluster ID using the following command: command line](#)
- [AGT0267E Failed to get the names of all the resource groups in the cluster.](#)
- [AGT0268E Failed to get the list of all the resource groups.](#)
- [AGT0269E Failed to get the service IP label information for resource group cluster resource group.](#)
- [AGT0270E Failed to get the volume group resources associated with cluster resource group cluster resource group.](#)
- [AGT0271I query command command selected.](#)
- [AGT0271E Failed to get the export resources associated with cluster resource group cluster resource group.](#)
- [AGT0272E Failed to get state of all the resource groups in the cluster.](#)
- [AGT0273E Failed to query the ODM query command.](#)
- [AGT0274E Storage Resource Agent initialization failed, return code: return code.](#)
- [AGT0275E Failed to get the HACMP node name using the following command:\ n command line](#)
- [AGT0276E Failed to get the physical volume information for the volume group volume group.](#)
- [AGT0277E Failed to get the logical volume information for the volume group volume group.](#)
- [AGT0278E Failed to get the SDD device information using the following command: Command](#)
- [AGT0279W The detected level of HACMP is not compatible. Please consult the user's guide for compatible versions.](#)
- [AGT0279I Registry entry added for installation location location.](#)
- [AGT0280W Failed to determine if the cluster is stable using the following command.](#)
- [AGT0281I Config file entry added for server server and port number number.](#)
- [AGT0281W The cluster is not ready to be probed. Waiting settle time seconds to retry \(retry count/retries\).](#)
- [AGT0282E Unable to probe the cluster.](#)
- [AGT0283I The cluster is ready and the probe will proceed.](#)
- [AGT0284E Agent host name does not accept scripts from server \(scripts are "Disabled"\). Script script name cannot run.](#)
- [AGT0285E The script script name has an incorrect Windows extension. The accepted Windows extensions are: extension list.](#)
- [AGT0301I Extracting file](#)
- [AGT0304E Failed to send agent registration for agent in server to server server.](#)
- [AGT0305I Successfully sent Agent, server, registration to server number](#)
- [AGT0306E Unable to send status to server.](#)
- [AGT0307I Successfully sent Probe complete status to server, server.](#)
- [AGT0314E Unable to stop all jobs for server server name](#)
- [AGT0383I Install completed successfully.](#)
- [AGT0389E Installation aborted. Communication method specified for this installation is communication method and does not match communication method, the communication method of the installed SRA.](#)
- [AGT0404I Creating Common Agent Package file file name](#)
- [AGT0405I Creating Common Agent Package directory directory name](#)
- [AGT0406I Common Agent Package file file name](#)
- [AGT0407I Common Agent Package directory directory name](#)
- [AGT0408E Failed to create Common Agent Package directory directory name](#)
- [AGT0409E Failed to create Common Agent Package file file name](#)

- [AGT0410E Could not delete Common Agent Package file file name](#)
- [AGT0411E SRAutil: invalid command command name](#)
- [AGT0412E SRAutil: invalid option options name](#)
- [AGT0413E Unable to start script script name](#)
- [AGT0414I Script script name](#)
- [AGT0415I Unable to load HBA library, rc: return code](#)
- [AGT0416I Number of HBA adapters on the system: number of adapters](#)
- [AGT0417I Adapter adapter name](#)
- [AGT0418I Found switched fabric: fabric ID](#)
- [AGT0419I Interconnected element element ID](#)
- [AGT0420I There are no adapters connected to switch fabric](#)
- [AGT0421E Unable to create fabric fabric component data file: data file](#)
- [AGT0422E Error writing to fabric fabric component data file: data file](#)
- [AGT0423I fabric command command on switch fabric failed, reason: reason code, explanation: explanation code](#)
- [AGT0424I Storage Resource Agent was unable to retrieve the fabric name for switch fabric fabric ID, reason: reason code, explanation: explanation code](#)
- [AGT0426I Fabric Probe Data: data field](#)
- [AGT0427I Fabric Discovery Data: data field](#)
- [AGT0428I Adapter adapter index not connected to switch fabric.](#)
- [AGT0429I Process process name invoked with command line arguments command arguments](#)
- [AGT0430I Process process name exiting with return code return code](#)
- [AGT0431W HBA API call function name failed with return code t](#)
- [AGT0432E Insufficient response buffer size passed for command command name](#)
- [AGT0433E response for command command name contains invalid data](#)
- [AGT0434E Errors parsing command name](#)
- [AGT0435I Command : command name](#)
- [AGT0436I Command : command arguments](#)
- [AGT0437I Error error code on adapter adapter index, unable to determine connection to switch fabric.](#)
- [AGT0438I Response: data field](#)
- [AGT0439E Errors writing command name](#)
- [AGT0440E Zone control command execution failed](#)
- [AGT0441E Unable to verify command completion, rc: return code](#)
- [AGT0442I ----- BEGIN OUTPUT -----](#)
- [AGT0443I ----- END OUTPUT -----](#)
- [AGT0444I Switched fabric: fabric ID already found](#)
- [AGT0445I Operation has been cancelled](#)
- [AGT0446I Fabric discovery successfully found switch fabric.](#)
- [AGT0447I Fabric discovery did not find switch fabric.](#)
- [AGT0448I Fabric discovery is already running.](#)
- [AGT0449I Fabric discovery failed, unable to discover switch fabric.](#)
- [AGT0452E No Multipath Device mapped to ID id](#)
- [AGT0453E Error when setting policy on device device: error](#)
- [AGT0454I For Multipath DM driver, only Round Robin policy is available.](#)
- [AGT0455W Setting policy for Multipath EMC Powerpath driver is not supported.](#)
- [AGT0456E No supported multipath driver was found on this system.](#)
- [AGT0457E Multipath policy configuration is not supported for this multipath driver.](#)
- [AGT0458I Setting multipath policy policy for following multipath devices: devices](#)
- [AGT0459I Executing Walk the Bus action for refreshing system configuration.](#)
- [AGT0460E Invalid source or target directory \(directory\). Source directory must be the installation image location. Target directory cannot be the same as source directory.](#)
- [AGT0461E Installation path contains an invalid character for the target platform: character](#)
- [AGT0462E At least one directory component in install location contains a reserved name for the target platform: name](#)
- [AGT0463I DM Multipath driver is installed but not loaded.](#)
- [AGT0464I multipath.conf could not be found on the system.](#)
- [AGT0469I Not enough disk space on disk. At least sizeMB is required!](#)
- [AGT0484E Cannot obtain the hostname of the Storage Resource Agent.](#)
- [AGT0485I The Storage Resource Agent on the IBM Spectrum Control server cannot be deleted.](#)
- [AGT0486W The available disk space is low on disk. It is recommended that at least size MB of disk space is available on the disk partition for successful operation of the Storage Resource agent.](#)
- [AGT0487E Not enough disk space available on disk. At least 20 MB of available disk space is required to complete the probe. Increase the available disk space on the disk partition and then start the probe again.](#)
- [AGT0504I Validating user user name](#)
- [AGT0505I User user name validation succeeded.](#)
- [AGT0506E User user name validation failed.](#)
- [AGT0507W User user name does not exist, user will be created.](#)
- [AGT0509I User user name created successfully.](#)
- [AGT0510E Failed to create user user name](#)
- [AGT0511E Usage error: -duser is missing.](#)
- [AGT0512E Usage error: -dpassword option is missing.](#)
- [AGT0513I Successfully sent Scan complete status to server, agent install directory](#)
- [AGT0514E Failed to send Scan complete status to server, agent install directory](#)
- [AGT0515I Stopping all SRA jobs ...](#)

AGT0001E Cannot set working directory to *directory*

Explanation

The specified working directory cannot be correctly accessed. If unable to determine the cause of the errors, contact IBM customer technical support.

Action

Check the working directory path and privileges.

Related reference

- [Getting support](#)

AGT0002E Invalid invocation: *swtchusr*.

Explanation

There was an invalid invocation of the swtchusr command.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0003I Agent registered.

Explanation

The agent has been registered.

AGT0004I Agent started.

Explanation

The agent has been started.

AGT0005I Waiting for swtchusr.

Explanation

Waiting for the swtchusr command to complete.

AGT0006E Process *process number* is already connected to the server.

Explanation

The identified process is already connected to the server.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0007E Another process is already connected to the server.

Explanation

Another process is already connected to the server, and only one process can be connected to the server at one time.

Action

Kill all processes and restart the server.

Related reference

- [Getting support](#)

AGT0008E Job-type should be *job type*, but it is *job type*.

Explanation

The job type for the scheduled job is invalid.

Action

Contact IBM customer technical support.

AGT0009E Error writing to swtchusr.

Explanation

Error during writing data to the swtchusr command.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0010E Class not found: *class name*.

Explanation

The class identified was not found in the Classpath.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0011E *class name* does not inherit from *class name*.

Explanation

The class identified is not inherited from the correct class.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0012E *class name* does not have appropriate constructor.

Explanation

The identified class does not have appropriate constructor.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0013E Error constructing *class name*.

Explanation

An error occurred trying to construct the identified class.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0014E Cannot close the log-file *log file name*

Explanation

The log-file cannot be properly closed resulting in possible lose of logging data.

Action

Ensure that the log file still exists and that is not being currently used by another process. If unable to determine the cause of the errors, contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0015E The process identified by the *process number* failed to start generating the *error status number*.

Explanation

The process failed to start properly.

Action

Check that the process is not hanging or blocked by another process. If unable to determine the cause of the errors, contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0016E Read from pipe failed.

Explanation

An error occurred while trying to read from a pipe.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0017E Write to pipe failed.

Explanation

An error occurred while trying to write to a pipe.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0018E Close of pipe failed.

Explanation

An error occurred while trying to close a pipe.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0019W Problem reading GUID identified by the *GUID number*.

Explanation

The GUID cannot be read.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0020E Unparseable class-loader string *class loader*.

Explanation

Error occurred while trying to obtain the binaries directory through the class loader parse.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0021E Error stopping the bundle.

Explanation

Error occurred while trying to stop a bundle.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0031E User *user name* does not exist.

Explanation

While trying to perform a swtchusr command, it was determined that the identified User does not exist. Verify that the User does exist.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0032E Cannot exec user command *shell shell -- errno = arguments error message*.

Explanation

An error occurred while trying to execute the specified command.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0033E *initgroups command(shell) fails -- errno = arguments error message.*

Explanation

The initgroups command for the User identified has failed with the reported error message.

Action

Make sure the user exists and has the proper authority, if the problem still persists contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0034E *Cannot create log-file in directory -- errno = arguments error message.*

Explanation

A log-file cannot be created in the specified directory.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0035W *Error reading logical volume.*

Explanation

The specified logical volume cannot be read.

AGT0036W *Error deserializing from file name.*

Explanation

An error has occurred trying to deserialize the job state from the specified file.

AGT0037W *Premature end of file -- file name.*

Explanation

An end of file condition occurred on the specified file before all information needed was read.

AGT0038W *Class class name not found restoring from file name.*

Explanation

The specified class was not found restoring from it from the specified file.

AGT0039W Object restored from *file name* is not *class name*, Object is *class name*.

Explanation

The object restored from specified file is not the class that was expected.

AGT0040E Agent Shutting down.

Explanation

The agent is shutting down, review the Agent logs for any error information.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0041E Agent not registered.

Explanation

The agent is not registered.

Action

The agent must be registered in Administrative Services.

AGT0042E Error writing *file name*.

Explanation

An error occurred while trying to write the specified filename.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0043E Error serializing to *file name*.

Explanation

An error has occurred trying to serialize an object to the specified filename.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0044E Fatal error -- cannot connect to self port number.

Explanation

An unrecoverable error has occurred while trying to connect to the specified port.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0045E Error parsing configuration file on line line number.

Explanation

An error has occurred while parsing the specified configuration file at the specified line number.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0046W In agent.config, key name is not an integer -- default of default key number used.

Explanation

In agent.config, the key name specified is not an integer value, the default value for that key will be used.

Action

Check the agent.config file for the specified Key Name and correct the value to be a valid integer. The agent must be restarted to pick up the change.

AGT0047E In agent.config, key name is not a valid port.

Explanation

In agent.config, the key name specified is not a valid port number.

Action

Check the agent.config file for the specified Key Name and correct the value to be a valid port number. The agent must be restarted to pick up the change.

AGT0048E In agent.config, no value supplied for key name.

Explanation

In agent.config, no value was supplied for the specified key name.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0049E Too many unreported jobs.

Explanation

Too many unreported jobs.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0050E Error parsing configuration file.

Explanation

An error has occurred while parsing the configuration file.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0051W The configuration file had at least one bad value.

Explanation

While parsing the configuration file, at least one bad value was found.

AGT0052W Cannot create temporary file in *directory*.

Explanation

A temporary file cannot be created in the specified directory.

AGT0053W Tried to send signal *process number* to nonexistent PID *process number*.

Explanation

Tried to send signal to nonexistent process.

AGT0054W No running job (*job name*, *job number*).

Explanation

A running job with the specified name and number were not found.

AGT0055E Jobs are not allowed to run as root.

Explanation

Jobs are not allowed to run as root.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0056I NOTICE: Server moved to *computer name:port number*.

Explanation

The server moved to the specified host and port number.

AGT0057E Unable to create message-queue from key *value*.

Explanation

Error occurred during the creation of the message queue.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0058E Error waiting for job to start.

Explanation

The time period for the starting of the specified job was exceeded.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0059I Received request to shut down *request type*.

Explanation

Received request to shut down agent.

AGT0060I Exiting normally.

Explanation

The agent is exiting normally.

AGT0061I Exiting abnormally.

Explanation

The agent has exited abnormally.

AGT0062E Cannot fork job.

Explanation

The attempt to fork the job failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0063E Error terminating process.

Explanation

A process terminated in an abnormal manner.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0064E Putative SID has too few hyphens: *security identifier*.

Explanation

The Putative SID has too few hyphens.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0065E Error waiting for shutdown request with *Queue ID*.

Explanation

The shutdown request was not processed in a normal manner.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

AGT0066E Error reaping.

Explanation

Error reaping.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

AGT0067E Error waiting for process *process name*.

Explanation

An error occurred while waiting for the specified process to complete.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

AGT0068E Cannot fork -- command: *command name*.

Explanation

An error has occurred trying to fork the specified command.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

AGT0069E open (/proc/uptime) fails.

Explanation

An error has occurred trying to obtain the last boot time.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0070E read(/proc/uptime) fails.

Explanation

A failure occurs trying to read the last boot time.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0071E /proc/uptime does not look right.

Explanation

The last boot time does not look right.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0072E Cannot determine system-boot time.

Explanation

The system-boot time cannot be determined.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0073E Problem creating up-time poller.

Explanation

A problem occurred while creating up-time poller.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0074E Unsupported Operating System.

Explanation

The agent is trying to be started on an unsupported Operating System.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0075E Cannot query *key name* *key value*.

Explanation

The registry key value specified cannot be queried.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0076E Performance title *data title* not found.

Explanation

The performance data title was not found in the registry.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0077E <<System>> performance object not found.

Explanation

The system performance data object was not found in the registry.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0078E The <<system>> performance counter not found.

Explanation

The system up time performance data counter was not found.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0079E Cannot create semaphore.

Explanation

A semaphore could not be created.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0080W Cannot retrieve process times.

Explanation

The time periods for the process execution could not be properly retrieved.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0081W Cannot retrieve process exit code.

Explanation

The process exit code could not be retrieved.

AGT0082E Cannot log in user *user name*.

Explanation

The specified user could not be logged in.

Action

Check that the user credentials are entered correctly and that it has appropriate access privileges. If unable to determine the cause of the errors, contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0083E Cannot obtain SID of local computer (*computer name*) .

Explanation

Cannot obtain SID of local computer.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0084E Cannot get name of local computer.

Explanation

The name of local computer cannot be retrieved.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0085E SID of local computer has unexpected type (*security identifier type*) .

Explanation

The Security Identifier of the local computer has an unexpected Security Identifier type.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0086W The GetTokenInformation method failed to execute.

Explanation

The token information could not be retrieved because the GetTokenInformation method failed to execute properly.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0087W The LookupAccountSid method failed to execute.

Explanation

The name of the account with the specified SID(security identifier type) could not be retrieved by the LookupAccountSid method.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0088W Cannot find any domain controller for domain *domain name*.

Explanation

The domain controller for the specified domain cannot be found.

Related reference

- [Getting support](#)

AGT0089W Cannot fetch the information for user *domain\user*.

Explanation

The user information could not be retrieved.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0090W Cannot load profile for user *username*.

Explanation

The user profile could not be loaded.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0091W Cannot unload user profile.

Explanation

The user profile could not be unloaded.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0092E Unable to retrieve status of job with the *job number*.

Explanation

The job status could not be retrieved.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0093E Cannot retrieve environment block.

Explanation

The system environment block for the users of the current process could not be retrieved.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0094E Cannot impersonate to create *resource*.

Explanation

The specified resource could not be created with the current credentials.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0095W Environment for user *username* not fully configured.

Explanation

The user environment is not fully configured.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0096E Malformed command-line.

Explanation

The command line is not correctly specified.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0097E Command not found: *fdisk* command.

Explanation

The fdisk command could not be found.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0098E Command has unsupported *extension*.

Explanation

The command does not support the specified file extension.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0099E Cannot duplicate handle.

Explanation

Cannot duplicate handle.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0101E Cannot determine whether this is a domain controller.

Explanation

It cannot be determined whether this is a domain controller.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0102E Cannot determine name of own domain.

Explanation

The name of the domain cannot be determined.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0103W Cannot make Windows Job Object.

Explanation

The Windows Job Object could not be created.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0104E Wrong server.

Explanation

The wrong server was communicated with.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0105W Ignoring server relocation in config file.

Explanation

If the server host and port were changed. The user is warned that they cannot do that.

AGT0106I Server has changed.

Explanation

The server has changed locations.

AGT0110E Cannot open key *key name* *key value*.

Explanation

The registry key HKEY_LOCAL_MACHINE\HARDWARE\DEVICEMAP\Scsi\Scsi Port cannot be opened.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0111I Rereading config file.

Explanation

The configuration file is being reread.

AGT0112E This product is not fully installed. To try again, stop and restart this agent.

Explanation

This product is not fully installed. To try again, stop and restart this agent.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0113E Cannot create temporary file in *directory* *error message*.

Explanation

A temporary file cannot be created in the specified agent home directory because of the error message specified.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0114E Cannot write to file *file name* *error message*.

Explanation

The specified file cannot be written because of the following error.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0115E Fork failed.

Explanation

The Fork command failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0116E Cannot exec *command name*.

Explanation

The command cannot be executed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0117W Cannot open *file* for auto-delete.

Explanation

The file cannot be opened for deletion.

Action

Check that there are no other processes which are accessing the file. If unable to determine the cause of the errors, contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0118E Upgrader is not okay.

Explanation

The agent upgrader failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0119I *version, modification and release.*

Explanation

An informational message that displays the Version Modification and Release of the product.

AGT0120E Error transmitting shutdown request to agent.

Explanation

An error occurred while transmitting shutdown request to the agent.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0121E Unable to create socket.

Explanation

An error occurred while trying to create a socket.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0122I The agent is already down.

Explanation

An informational message letting the user know that the agent is already down.

AGT0123W Unable to determine if agent is active. Sending shutdown request.

Explanation

It cannot be determined if the agent is active, so a shutdown request is being sent.

AGT0124E Unable to create restartable job directory: *directory name*.

Explanation

The restartable job directory cannot be created.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0125E Unable to read directory *directory name*.

Explanation

The contents of the directory cannot be read.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0126E Unable to restart job from file *file name*.

Explanation

The job from the specified file is unable to be restarted.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0127E Unable to restart job *job name.name* run number *run number*.

Explanation

The specified job with the specified run number is unable to be restarted.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0128W Cannot retrieve global structure.

Explanation

The global structure cannot be retrieved.

AGT0129W A system call failed in the agent program.

Explanation

A system call to determine the number of processors failed when called in the agent program.

AGT0130E The upgrader program parameters are missing.

Explanation

The parameters that are to be sent to the upgrader program are missing.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0131I Exit Status = *exit message*.

Explanation

This is an informational message that displays the exit status of the Agent.

AGT0132E Cannot find/create script *<script name>*.

Explanation

The specified script cannot be found/created.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0133I Running Command: *script name script contents*.

Explanation

This is an informational message letting you know what script is being run.

AGT0134E Putative SID does not start correctly: *security identifier*.

Explanation

The Putative Security Identifier specified did not start correctly.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0135E Component *component name* is not a number: *component number*.

Explanation

The Component specified is not a component number.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0136E No INSTANCEn.DAT file found!

Explanation

The INSTANCEn.DAT file could not be found. This file is needed to determine the location of Java.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0137E Cannot determine working directory!

Explanation

The working directory cannot be determined.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0138W *product name* is not licensed on this computer.

Explanation

The specified product is not licensed on this computer.

AGT0139E An initial probe cannot be performed.

Explanation

Because the product is not licensed, an initial probe cannot be performed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0140W Discovery will not be performed.

Explanation

The discovery process will not be performed.

AGT0141W A filesystem scan will not be performed.

Explanation

Because the product is not licensed, a filesystem scan will not be performed.

AGT0142E Cannot make handle inheritable.

Explanation

The stream handle cannot be made inheritable.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0143I *file_count* files scanned

Explanation

This message indicates the progress of the file system scan.

AGT0144I *file_count* total files scanned

Explanation

This message states the total number of files that were encountered during the file system scan.

AGT0145I Retrieving job definition from server

Explanation

Most agent jobs, as part of their initial processing, contact the server to retrieve job-control parameters. This message indicates that the agent is about to take that step.

AGT0146I Scan started

Explanation

The agent has begun scanning

AGT0147I Retrieving report partition *partition_number*

Explanation

Owing to their possibly large size, certain reports are retrieved piecemeal. This message, logged during a batch report, indicates that the agent is about to request the server for a segment of the report.

AGT0148I Report data retrieved

Explanation

All batch report data have been retrieved from the server. The agent will now write the data to their specified destination.

AGT0149I Retrieving history data

Explanation

During a batch report, the agent is about to request history data from the server.

AGT0150I Deleting temporary files

Explanation

The agent is about to request the server to clean up temporary files created during the report retrieval process.

AGT0151E Unable to retrieve report definition

Explanation

The agent could not retrieve the report job definition from the server. Common reasons for this problem include network errors, configuration errors, and repository SQL errors.

Action

Check the server's TPCD_nnnnnn.log and/or server_nnnnnn.log for further details. If these do not reveal sufficient information, contact IBM support.

Related reference

-  [Getting support](#)

AGT0152I Job definition retrieved

Explanation

An agent job has successfully retrieved its control parameters from the server.

AGT0153E MSCSEventListener thread interrupted. Shutting down native MSCS event listener.

Explanation

The Microsoft Cluster Service event listener thread was interrupted and is now shutting down.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0154E No IP address could be found for the local host.

Explanation

The IP address for the local host could not be determined.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0155I Windows Scan Option : *scan parameters*.

Explanation

Presents the current selected scan parameters for Windows systems.

AGT0156W Error killing process *process id*.

Explanation

The process with the specified process id could not be terminated.

Action

Check that the process is alive and is not blocked by another process. If unable to determine the cause of the errors, contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0157E Interrupted (agent shutting down) .

Explanation

The agent thread was interrupted and the agent will shut down.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

AGT0159E A problem was encountered stopping IBM Spectrum Control subagent.

Explanation

The shutdown of the data subagent did not complete successfully.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

AGT0160E Error enumerating keys under *registry path*.

Explanation

A error was encountered during the enumeration of the registry keys for the specified registry path.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

AGT0161W Error querying value *root path\\subkey name\\library file*.

Explanation

A error has occurred while trying to query the library file value with the specified registry path.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

AGT0162E *Library* : HBA_RegisterLibrary returned *value*.

Explanation

The HBA_RegisterLibrary method failed with the specified return code.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0163E Library : HBA_LoadLibrary returned value.

Explanation

The HBA_LoadLibrary method failed with the specified return code.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0164W Library : HBA_GetAdapterName(adapter number) returned adapter name.

Explanation

The HBA_GetAdapterName method failed for the specified adapter.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0165W Library : HBA_OpenAdapter(adapter name) failed.

Explanation

The HBA_OpenAdapter method failed for the specified adapter.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0166W The HBA API HBA_GetAdapterAttributes for adapter adapter name returned error error code.

Explanation

The HBA_GetAdapterAttributes API call failed with the indicated return code. Some of the standard codes are as follows: 1 - General Error 2 - Function is not supported 9 - SCSI check condition 10 - HBA is busy or reserved 11 - Request timed out 15 - Incompatible library and driver modules 16 - Multiple adapters have the same WWN 30 - Target busy.

Action

Contact your HBA vendor for support. If the behavior of the adapter with respect to the HBA_GetAdapterAttributes API is intended by the vendor, then ignore this warning message.

AGT0167W *Library : HBA_GetAdapterPortAttributes(adapter number, port) returned status.*

Explanation

The HBA_GetAdapterPortAttributes method failed for the adapter port with the specified status.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0168W *HBA data cannot be collected because the adapter name> adapter on the server does not support the HBA_GetAdapterAttributes function.*

Explanation

When a probe attempts to collect data about a server, this function call to the HBA API is not supported by the adapter on that server. Information about the Fibre Channel Storage Area Network that is typically provided by this function is not collected.

Action

For information about adapter support of the common HBA API, see the documentation that is provided by the vendor. If you require additional information, contact the HBA vendor. If the behavior of the adapter for the HBA_GetAdapterAttributes API call is intended by the vendor, you can ignore this warning message.

AGT0200E *Error waiting for process.*

Explanation

A error has occurred while waiting for a windows script to run.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0250E *Error starting bundle: exception message.*

Explanation

A error with the specified message has occurred while trying to start the bundle.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0251E *Failed to install language pack.*

Explanation

The scheduled maintenance job failed to install the language pack.

Action

Retry the scheduled maintenance job.

AGT0252E Failed to install language pack: File not found *archive*

Explanation

The scheduled maintenance job failed to install the language pack. The archive containing the language pack could not be found.

Action

Retry the scheduled maintenance job.

AGT0253E Invalid file format: *file name*

Explanation

The scheduled maintenance job failed to install the language pack. The specified file is not in the expected format.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0254E Failed to load language pack.

Explanation

The agent was not able to load the language pack. The agent may use the default locale.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0256I Waiting for Common Agent services.

Explanation

The agent startup will be performed when the required services have started.

Action

Make sure the Connector service of the Common Agent has started successfully. The service may fail to start if the credentials required for SSL have expired and the Common Agent has not renewed the credentials with the Agent Manager.

AGT0259E Agent cannot upgrade, the required space of 50 meg was not met.

Explanation

The agent upgrade process cannot be completed because there is not enough space on disk for the process to complete.

Action

Check that the upgrade location has at least 50 megabytes of free space for the agent upgrade. If unable to determine the cause of the errors, contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0260E Failed to create a session with Reliable Scalable Cluster Technology *Error ID-Error type: Error Description*

Explanation

IBM Spectrum Control attempted to make a connection with Reliable Scalable Cluster Technology (RSCT) to collect information about an HACMP cluster but the connection failed.

Action

Check the status of the RSCT subsystems. If the problem continues contact IBM customer support.

Related reference

- [Getting support](#)

AGT0261E Failed to end a session with Reliable Scalable Cluster Technology *Error ID-Error type: Error Description*

Explanation

IBM Spectrum Control failed to end a connection with Reliable Scalable Cluster Technology (RSCT) that was used to collect information about an HACMP cluster.

Action

Check the status of the RSCT subsystems. If the problem continues contact IBM customer support.

Related reference

- [Getting support](#)

AGT0262W Concurrent cluster resource groups are not supported. The cluster will not be probed.

Explanation

The computer being probed is a member of an operating system cluster that has at least one concurrent resource group defined. Since IBM Spectrum Control does not support concurrent resource groups the probe will not collect information about the cluster.

Action

Remove the concurrent resource group from the cluster and retry the probe.

AGT0263E Failed to get resource data from the HACMP cluster using the RSCT RMC interface.

Explanation

The attempt to query the RSCT RMC subsystem for cluster resource information failed.

Action

Check the status of the RSCT subsystems. If the problem continues contact IBM customer support.

Related reference

- [Getting support](#)

AGT0264E Failed to determine if the local node is clustered using the following command: *command line*

Explanation

Failed to execute or parse the output of the specified command.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0265E Failed to get the cluster name using the following command: *command line*

Explanation

Failed to execute or parse the output of the specified command.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0266E Failed to get the cluster ID using the following command: *command line*

Explanation

Failed to execute or parse the output of the specified command.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0267E Failed to get the names of all the resource groups in the cluster.

Explanation

The attempt to query the RSCT RMC subsystem for the list of cluster resource groups failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0268E Failed to get the list of all the resource groups.

Explanation

The attempt to query the RSCT RMC subsystem for the list of cluster resource groups failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0269E Failed to get the service IP label information for resource group *cluster resource group*.

Explanation

The attempt to query the RSCT RMC subsystem for the list of service IP labels associated with the specified cluster resource group failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0270E Failed to get the volume group resources associated with cluster resource group *cluster resource group*.

Explanation

The attempt to query the RSCT RMC subsystem for the list of volume groups associated with the specified cluster resource group failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0271I *query command* command selected.

Explanation

The command or commands listed were selected for execution.

Action

No response is necessary. This message is for information only.

AGT0271E Failed to get the export resources associated with cluster resource group *cluster resource group*.

Explanation

The attempt to query the RSCT RMC subsystem for the list of exports associated with the specified cluster resource group failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0272E Failed to get state of all the resource groups in the cluster.

Explanation

The attempt to query the RSCT RMC subsystem for the state of the cluster resource groups failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0273E Failed to query the ODM *query command*.

Explanation

The specified query command for the Object Data Manager(ODM) was unsuccessful.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0274E Storage Resource Agent initialization failed, return code: *return code*.

Explanation

Storage Resource Agent initialization failed with the indicated return code, the Storage Resource Agent cannot proceed.

Action

Refer to the Storage Resource Agent return code documentation. When performing a Storage Resource Agent install, verify the specified installation directory is writable. If the directory is not empty, choose another directory or use the -force option to proceed with the installation.

Related reference

- [Getting support](#)

AGT0275E Failed to get the HACMP node name using the following command:\ \n *command line*

Explanation

Failed to execute or parse the output of the specified command.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0276E Failed to get the physical volume information for the volume group *volume group*.

Explanation

Failed to get the list of physical volumes associated with the specified volume group.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0277E Failed to get the logical volume information for the volume group *volume group*.

Explanation

Failed to get the list of logical volumes associated with the specified volume group.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0278E Failed to get the SDD device information using the following command: *Command*

Explanation

Failed to execute or parse the output of the specified command.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0279W The detected level of HACMP is not compatible. Please consult the user's guide for compatible versions.

Explanation

The detected level of HACMP is not compatible. Please consult the user's guide for compatible versions. The probe will continue but may experience errors.

AGT0279I Registry entry added for installation location *location*.

Explanation

A registry entry was added for the location indicated.

Action

No response is necessary. This message is for information only.

AGT0280W Failed to determine if the cluster is stable using the following *command*.

Explanation

The HACMP cluster specified command could not determine if the cluster is stable.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0281I Config file entry added for server *server* and port number *number*.

Explanation

An entry for this server and port were added to the config file.

Action

No response is necessary. This message is for information only.

AGT0281W The cluster is not ready to be probed. Waiting *settle time* seconds to retry (*retry count/retries*).

Explanation

The cluster is not ready to be probed and will try a new attempt at a later time based on the left number of retries.

Action

Verify the status of Cluster Manager daemon (clstrmgrES) using the following command:

```
/usr/bin/lsrsc -ls clstrmgrES
```

If the output of the command indicates that the daemon is not running or the 'Current State:' line does not contain ST_INIT or ST_STABLE status, then start the Cluster Manager daemon using the following command:

```
startsrc -s clstrmgrES
```

AGT0282E Unable to probe the cluster.

Explanation

The cluster could not be probed after repeated attempts.

Action

Verify the status of Cluster Manager daemon (clstrmgrES) using the following command:

```
/usr/bin/lsrsc -ls clstrmgrES
```

If the output of the command indicates that the daemon is not running or the 'Current State:' line does not contain ST_INIT or ST_STABLE status, then start the Cluster Manager daemon using the following command:

```
startsrc -s clstrmgrES
```

AGT0283I The cluster is ready and the probe will proceed.

Explanation

The probe initialization process completed with success.

AGT0284E Agent *host name* does not accept scripts from server (scripts are '*Disabled*'). Script *script name* cannot run.

Explanation

The specified script cannot be run because the agent scripts are disabled.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0285E The script *script name* has an incorrect Windows extension. The accepted Windows extensions are: *extension list*.

Explanation

The specified script has a extension that is not one of the supported Windows script extensions.

Action

Correct the script extension so that it matches the supported Windows script extensions. If unable to determine the cause of the errors, contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0301I Extracting *file*

Explanation

The agent process is extracting files.

Action

No response is necessary. This message is for information only.

AGT0304E Failed to send agent registration for agent in *server* to *server server*.

Explanation

The agent was unable to send registration information to the IBM Spectrum Control server

Action

Check logs 'server_XXXXXX.log' and 'TPCD_XXXXXX.log' on server in directory IBM Spectrum Control Server install dir/data/log'. Validate agent hostname on server side. Authenticate with the firewall if necessary.

AGT0305I Successfully sent Agent, *server*, registration to *server number*

Explanation

Agent registration has been sent to the server successfully.

Action

No response is necessary. This message is for information only.

AGT0306E Unable to send status to *server*.

Explanation

The agent could not send the probe results to the server.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0307I Successfully sent Probe complete status to *server*, *server*.

Explanation

Probe complete status has been sent to the server.

Action

No response is necessary. This message is for information only.

AGT0314E Unable to stop all jobs for server *server name*

Explanation

Storage Resource Agent was unable to stop all running jobs for the specified server

Action

A manual stop of running jobs may be required. Use operating system commands to identify and to stop running Storage Resource Agent jobs.

AGT0383I Install completed successfully.

Explanation

The agent install completed successfully.

Action

This message is informational. No action is required.

AGT0389E Installation aborted. Communication method specified for this installation is *communication method* and does not match *communication method*, the communication method of the installed SRA.

Explanation

Changing the existing communication method from daemon to non-daemon or vice versa is not supported.

Action

If you want to change the agent communication method, uninstall the agent and reinstall using the new communication method.

AGT0404I Creating Common Agent Package file *file name*

Explanation

Common Agent Package file is created

AGT0405I Creating Common Agent Package directory *directory name*

Explanation

Common Agent Package directory is created

AGT0406I Common Agent Package file *file name*

Explanation

Common Agent Package file name

AGT0407I Common Agent Package directory *directory name*

Explanation

The Common Agent Package directory name

AGT0408E Failed to create Common Agent Package directory *directory name*

Explanation

Error creating Common Agent Package directory where the Common Agent Package file is stored

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0409E Failed to create Common Agent Package file *file name*

Explanation

Error creating Common Agent Package file

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0410E Could not delete Common Agent Package file *file name*

Explanation

Error deleting Common Agent Package file

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0411E SRAutil: invalid command *command name*

Explanation

Invalid command passed to SRAutil component

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0412E SRAutil: invalid option *options name*

Explanation

Invalid option passed to SRAutil component

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0413E Unable to start script *script name*

Explanation

Error starting script

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0414I Script *script name*

Explanation

Script name to execute

AGT0415I Unable to load HBA library, rc: *return code*

Explanation

HBA API library could not be loaded. The system may not have any HBA Adapters configured.

AGT0416I Number of HBA adapters on the system: *number of adapters*

Explanation

Numbe of adapters discovered on the system

AGT0417I Adapter *adapter name*

Explanation

Name of HBA adapter on the system

AGT0418I Found switched fabric: *fabric ID*

Explanation

Discovered switched fabric

AGT0419I Interconnected element *element ID*

Explanation

Interconnected elements in switched fabric

AGT0420I There are no adapters connected to switch fabric

Explanation

Adapters are not connected to switch fabric

AGT0421E Unable to create fabric *fabric component* data file: *data file*

Explanation

Storage Resource Agent fabric component was unable to create the output data file

Action

Verify the Storage Resource Agent installation directory disk space is not full

AGT0422E Error writing to fabric *fabric component* data file: *data file*

Explanation

Storage Resource Agent fabric component was unable to write to the output data file.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0423I *fabric command* command on switch fabric failed, reason: *reason code*, explanation: *explanation code*

Explanation

HBA command on switch fabric failed

AGT0424I Storage Resource Agent was unable to retrieve the fabric name for switch fabric *fabric ID*, reason: *reason code*, explanation: *explanation code*

Explanation

The GFN command on switch fabric failed

AGT0426I Fabric Probe Data: *data field*

Explanation

fabric probe data

AGT0427I Fabric Discovery Data: *data field*

Explanation

fabric discovery data

AGT0428I Adapter *adapter index* not connected to switch fabric.

Explanation

The adapter referenced by the specified index is not connected to switch fabric.

Action

If adapter is connected to switch fabric, inspect the Storage Resource Agent Fabric Discovery log file for any errors.

AGT0429I Process *process name* invoked with command line arguments *command arguments*

Explanation

Process name and arguments of the Storage Resource Agent process being invoked

AGT0430I Process *process name* exiting with return code *return code*

Explanation

Exit code of the Storage Resource Agent process

AGT0431W HBA API call *function name* failed with return code *t*

Explanation

The HBA API function failed

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0432E Insufficient response buffer size passed for command *command name*

Explanation

Fabric command failed due to insufficient response buffer size

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0433E response for command *command name* contains invalid data

Explanation

Fabric command response data was invalid

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0434E Errors parsing *command name*

Explanation

Storage Resource Agent fabric zone control component was unable to parse a zone control command

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0435I Command : *command name*

Explanation

Fabric zone control command being processed

AGT0436I Command : *command arguments*

Explanation

Fabric zone control command arguments

AGT0437I Error *error code* on adapter *adapter index*, unable to determine connection to switch fabric.

Explanation

The Storage Resource Agent Fabric Discovery encountered an error when trying to determine if adapter referenced by the specified index is connected to a switch fabric.

Action

Inspect the Storage Resource Agent Fabric Discovery log file for errors.

AGT0438I Response: *data field*

Explanation

Fabric zone control response data

AGT0439E Errors writing *command name*

Explanation

Fabric Zone Control component was unable to write to the output data file

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0440E Zone control command execution failed

Explanation

An error was encountered processing fabric zone control commands

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0441E Unable to verify command completion, rc: *return code*

Explanation

An error was encountered while waiting for a command completion.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0442I ----- BEGIN OUTPUT -----

Explanation

Command output begin

AGT0443I ----- END OUTPUT -----

Explanation

Command output end

AGT0444I Switched fabric: *fabric ID* already found

Explanation

The fabric has already been discovered

AGT0445I Operation has been cancelled

Explanation

Operation was cancelled

AGT0446I Fabric discovery successfully found switch fabric.

Explanation

The fabric discovery process successfully found the switch fabric connected to the system

AGT0447I Fabric discovery did not find switch fabric.

Explanation

The fabric discovery process did not find switch fabric.

AGT0448I Fabric discovery is already running.

Explanation

The fabric discovery process was not performed because a fabric discovery process is already running.

AGT0449I Fabric discovery failed, unable to discover switch fabric.

Explanation

The fabric discovery process failed and was unable to determine if system is connected to switch fabric.

Action

The commands used to discover the switch fabric failed. Inspect previous messages to determine the command failure.

AGT0452E No Multipath Device mapped to ID *id*

Explanation

No corresponding SDD device was found for specified volume serial.

Action

Make sure that the volume with specified ID has been assigned and it is visible to the host.

AGT0453E Error when setting policy on device *device*: error

Explanation

An error was encountered when setting the multipath policy for specific device id.

Action

Refer to the agent log file for more details.

AGT0454I For Multipath DM driver, only Round Robin policy is available.

Explanation

For Multipath DM driver, only Round Robin policy is available.

Action

If different multipath policy is required, then a supported multipath driver that supports this policy is required.

AGT0455W Setting policy for Multipath EMC Powerpath driver is not supported.

Explanation

Multipath policy configuration is not supported for Multipath EMC Powerpath driver.

Action

Make sure that the multipath driver supports the requested multipath policy.

AGT0456E No supported multipath driver was found on this system.

Explanation

No supported multipath driver is installed on this system.

Action

Make sure that fully supported multipath driver is available on this system for Path Planner configurations.

AGT0457E Multipath policy configuration is not supported for this multipath driver.

Explanation

Multipath policy configuration is not supported for this multipath driver.

Action

Make sure that the multipath driver supports the requested multipath policy.

AGT0458I Setting *multipath policy* policy for following multipath devices: *devices*

Explanation

Setting the specified policy type for the specified multipath devices list.

Action

No action is required.

AGT0459I Executing Walk the Bus action for refreshing system configuration.

Explanation

The system is rescanned for renewing the configuration.

Action

No action is required.

AGT0460E Invalid source or target directory (*directory*). Source directory must be the installation image location. Target directory cannot be the same as source directory.

Explanation

The source and target directory cannot be the same.

Action

Please make sure that the Agent executable from the installation image location is being invoked and that the installation location specified for the install points to a different location than the one from where installation is being invoked.

AGT0461E Installation path contains an invalid character for the target platform: *character*

Explanation

The character *character* cannot be part of the installation path.

Action

Please choose an install directory which does not contain reserved or shell special characters on the platform where the installation is attempted.

AGT0462E At least one directory component in install location contains a reserved name for the target platform: *name*

Explanation

The name *name* cannot be part of the installation path.

Action

Please choose an install directory which does not contain reserved names on the platform where the installation is attempted.

AGT0463I DM Multipath driver is installed but not loaded.

Explanation

Ensure that the Device Mapper packages are installed on the system and available. Ensure the latest HBA drivers are installed.

Action

Verify whether the dm-multipath modules are loaded and the Multipathd Daemon has started automatically at boot time.

AGT0464I multipath.conf could not be found on the system.

Explanation

/etc/multipath.conf configuration file was not found on the system or it is not configured.

Action

The configuration can be done based on a sample multipath.conf configuration file. You can run /sbin/mpathconf to create or modify /etc/multipath.conf.

AGT0469I Not enough disk space on disk. At least sizeMB is required!

Explanation

Minimum disk space is required for the Storage Resource Agent on the IBM Spectrum Control server.

Action

Refer to the IBM Spectrum Control documentation.

AGT0484E Cannot obtain the hostname of the Storage Resource Agent.

Explanation

The hostname of Storage Resource Agent could not be determined.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0485I The Storage Resource Agent on the IBM Spectrum Control server cannot be deleted.

Explanation

The Storage Resource Agent on the IBM Spectrum Control server is a required component.

Action

Refer to the IBM Spectrum Control documentation.

AGT0486W The available disk space is low on *disk*. It is recommended that at least *size* MB of disk space is available on the disk partition for successful operation of the Storage Resource agent.

Explanation

To enable the Storage Resource agent to operate efficiently, a minimum amount of available disk space is required. Check the product support matrix for information about the recommended disk space for Storage Resource agents.

Action

Increase the available disk space to the recommended values in the product support matrix.

AGT0487E Not enough disk space available on *disk*. At least 20 MB of available disk space is required to complete the probe. Increase the available disk space on the disk partition and then start the probe again.

Explanation

A probe cannot complete successfully unless at least 20 MB space is available on the disk partition. This space is required to write the probe logs.

Action

No additional action is required.

AGT0504I Validating user *user name*

Explanation

User name validation starts.

AGT0505I User *user name* validation succeeded.

Explanation

User name was validated.

AGT0506E User *user name* validation failed.

Explanation

User name validation failed

Action

Check that a valid user name is specified.

AGT0507W User *user name* does not exist, user will be created.

Explanation

User name does not exist, the user is created

Action

None

AGT0509I *User user name created successfully.*

Explanation

User was created successfully

AGT0510E *Failed to create user user name*

Explanation

Unable to create user.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0511E *Usage error: -duser is missing.*

Explanation

Expected user option is missing.

Action

Specify a user with the -duser option

AGT0512E *Usage error: -dpassword option is missing.*

Explanation

Expected password option is missing.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

AGT0513I *Successfully sent Scan complete status to server, agent install directory*

Explanation

Storage Resource Agent successfully sent the File System Scan job completion status to the server

AGT0514E Failed to send Scan complete status to server, agent install directory

Explanation

Storage Resource Agent could not send the File System Scan job completion status to the server

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

AGT0515I Stopping all SRA jobs ...

Explanation

Storage Resource Agent is stopping all the running jobs.

ALR - Spectrum Control Alert messages

- [ALR0001I The amount of RAM on host server name has changed from current value to new value.](#)
- [ALR0002I The amount of virtual memory on host server name has changed from current value to new value.](#)
- [ALR0003I A new disk drive has been detected on host server name. Disk manufacturer/serial Number: manufacturer/serial Number.](#)
- [ALR0004E A previously visible disk drive can no longer be found on host server name. Disk manufacturer/serial number: manufacturer/serial number.](#)
- [ALR0005I A new filesystem has been detected on host server name. Filesystem mount point: mount point.](#)
- [ALR0006E A previously visible filesystem can no longer be found on host server name. Filesystem mount point: mount point.](#)
- [ALR0007E A disk drive visible on host server name has predicted that a disk failure is imminent. Disk manufacturer/serial number: manufacturer/serial number.](#)
- [ALR0008I The physical space definition of filesystem filesystem name on host server name has been reconfigured.](#)
- [ALR0009W The free space on filesystem filesystem name on host server name has fallen below the threshold value of threshold. The free space is freespace or percent of the filesystem capacity.](#)
- [ALR0010W The number of free inodes on filesystem filesystem name on host server name has fallen below the threshold value of threshold. The number of free inodes is free inodes or percent of the filesystem's total inodes.](#)
- [ALR0011W A new grown defect has been detected on a disk visible to host server name. Disk manufacturer/serial Number: disk manufacturer/serial number. Current grown defects: current, Previous grown defects: previous.](#)
- [ALR0012W The number of grown disk defects has exceeded the threshold value of threshold. Host: server name, Disk manufacturer/serial number: disk manufacturer/serial number, Current grown defects: current, Previous grown defects: previous.](#)
- [ALR0013W A new monitored directory has been detected on host server name. Directory name: directory name, Directory Group: directory group.](#)
- [ALR0014E A monitored directory has been removed from host server name. Directory name: directory name, Directory group: directory group.](#)
- [ALR0015W Directory directory on host server name has exceeded its space usage quota of quota. The directory is currently consuming usage or percent of the filesystem capacity.](#)
- [ALR0016W Filesystem filesystem name on host server name has violated a filesystem constraint. number of files file\(s\) consuming space or percent of the filesystem capacity are in violation of the conditions defined in this constraint. The constraint threshold is threshold. User user name} has number of files} files consuming space} of storage. violating owners}.](#)
- [ALR0017E Host server name appears to be down. number of attempts attempt\(s\) to ping this host have failed.](#)
- [ALR0018W quota name user has exceeded a network storage usage quota of usage. This user is currently consuming amount of storage.](#)
- [ALR0019W quota name user on host server name has exceeded a server storage usage quota of usage. This user is currently consuming amount of storage.](#)
- [ALR0020W quota name user on host server name has exceeded a filesystem usage quota of <usage>valueon filesystem filesystem name. This user is currently consuming amount of storage.](#)
- [ALR0021W Run number number of job creator job name has failed on run number of total jobs total jobs.](#)
- [ALR0022I Server server name has been discovered.](#)
- [ALR0023W Run number number of job creator job name has failed.](#)
- [ALR0024W User user name on host server name has exceeded a table space usage quota of usage on table space rdbms type, table space instance. This user is currently consuming amount of storage.](#)
- [ALR0025W User user name on host server name has exceeded an RDBMS instance usage quota of quota on rdbms type instance instance name. This user is currently consuming amount of storage.](#)
- [ALR0026W User user name has exceeded a network database storage usage quota of quota. This user is currently consuming amount of storage.](#)
- [ALR0027W The log directory file name archived log directory value on host server name has exceeded the threshold value of threshold. This directory currently contains number of logs archived logs consuming amount of storage.](#)
- [ALR0028I A new table space name has been discovered on rdbms type instance on host server name. rdbms instance: database.](#)
- [ALR0029E value value has been dropped. RDBMS: value value, Host: host name.](#)
- [ALR0030E value value has been taken offline. RDBMS: value value, Host: host name.](#)
- [ALR0031W The free space on value: value, RDBMS: value value, host: host name, has fallen below the threshold value of value. The free space is value or value of the value capacity.](#)
- [ALR0032W The free space on table space: table space name, RDBMS: rdbms type instance, database: database, host: server name, is fragmented across number of extents extents. This exceeds the threshold value of threshold extents. The largest contiguous free extent is largest extent.](#)

- [ALR0033W](#) The largest free extent available on table space: table space name, RDBMS: rdbms type rdbms name, database: database name, host: server name, has fallen below the threshold value of threshold. The largest free extent is largest extent.
- [ALR0034W](#) Segment segment name of table/cluster table name on host: server name, RDBMS: rdbms type rdbms name, database: database name, is fragmented across number of extents extents. This exceeds the threshold value of threshold extents. This segment is a value type segment.
- [ALR0035W](#) Segment value of table/cluster value on host: host name, RDBMS: value value, database: value, is nearing the maximum number of extents available to it. This segment currently occupies value extents. The value additional extent(s) available to this segment falls below the defined threshold of value extent(s). The segment is a type type segment.
- [ALR0036W](#) Table/cluster table name on host: server name, RDBMS: rdbms type rdbms name, database: database name, has exceeded a space usage quota of usage quota. This table is currently consuming amount of storage.
- [ALR0037W](#) Table/cluster value on host: host name, RDBMS: value value, database: database name, has exceeded a chained row quota of value. Statistics indicate that value rows or value of the total rows are chained.
- [ALR0038W](#) Segment name of table/cluster table name on host: server name, RDBMS: rdbms type rdbms name, database: database name, has amount of unused, wasted space. This represents amount of the total space allocated to the segment, and exceeds the threshold value of threshold. This segment is a type type segment.
- [ALR0039E](#) Table/cluster table name on host: server name, RDBMS: rdbms type rdbms name, database: database name, has been dropped.
- [ALR0040W](#) Filer filer name has been discovered. Spectrum Control will not monitor this filer until it has been licensed.
- [ALR0041W](#) The amount of log freespace available on instance: database, RDBMS: rdbms type, <rdbms>host: server name, has fallen below the threshold value of threshold. The amount of log freespace available is amount or percent of the total capacity.
- [ALR0042I](#) A new device has been discovered on rdbms instance instance on host server name. Device: device name, Capacity: capacity, File Name: filename.
- [ALR0043E](#) Device device name has been dropped from rdbms instance instance on host server name. Capacity: capacity, File Name: filename.
- [ALR0044W](#) The amount of device freespace available on value instance value on host value has fallen below the threshold value of value. The amount of device freespace available is value or value of the current capacity of value.
- [ALR0045W](#) The amount of device freespace available on device instance instance on host server name has gone above the threshold value of threshold. The amount of device freespace available is freespace or percent of the current capacity of total capacity.
- [ALR0046W](#) Database database name has not been backed up in the last number of days days. Last backup for the database was on date. RDBMS: rdbms type rdbms name, Host: server name.
- [ALR0047W](#) filer name filer type is no longer accessible from host entity.
- [ALR0048W](#) Storage Subsystem subsystem name is no longer accessible from host server name.
- [ALR0049W](#) disk array name disk array type has been discovered from host server name. Spectrum Control will not monitor this disk array until it has been selected as for monitoring from within the Storage Subsystem Administration GUI.
- [ALR0050W](#) The amount of cache on storage subsystem subsystem name has changed from old value to new value.
- [ALR0052W](#) Filesystem filesystem name on host server name will be automatically extended because its free space has fallen below the threshold of threshold. Current free space: current free space; Current capacity: current capacity; Target capacity: target capacity).
- [ALR0053W](#) Filesystem filesystem name on host server name needs extension but will not be because its current capacity of current capacity exceeds the specified limit of limit. Filesystem free space: freespace (current capacity of current capacity).
- [ALR0055I](#) Cluster resource group cluster resource group name was added to cluster cluster name on node node name.
- [ALR0056I](#) Cluster resource group cluster resource group name was removed from cluster cluster name on node node name.
- [ALR0057I](#) Cluster resource group cluster resource group name was moved in cluster cluster name from node node name to node node name.
- [ALR0076W](#) Performance monitor failure for device value.
- [ALR0500E](#) The Disk Utilization Percentage of array array name in storage system storage system name was measured to be measured value%. This violated the critical-stress boundary value of boundary value%.
- [ALR0501W](#) The Disk Utilization Percentage of array array name in storage system storage system name was measured to be measured value%. This violated the warning-stress boundary value of boundary value%.
- [ALR0502W](#) The Disk Utilization Percentage of array array name in storage system storage system name was measured to be measured value%. This violated the warning-idle boundary value of boundary value%.
- [ALR0503E](#) The Disk Utilization Percentage of array array name in storage system storage system name was measured to be measured value%. This violated the critical-idle boundary value of boundary value%.
- [ALR0504E](#) The Total Back-end I/O Rate of array, MDisk, or MDisk Group Array, MDisk, or MDisk Group name in storage system storage system name was measured to be measured value ops/s, which violated the defined critical-stress boundary value of boundary value ops/s.
- [ALR0505W](#) The Total Back-end I/O Rate of array, MDisk, or MDisk Group Array, MDisk, or MDisk Group name in storage system storage system name was measured to be measured value ops/s, which violated the defined warning-stress boundary value of boundary value ops/s.
- [ALR0506W](#) The Total Back-end I/O Rate of array, MDisk, or MDisk Group Array, MDisk, or MDisk Group name in storage system storage system name was measured to be measured value ops/s, which violated the defined warning-idle boundary value of boundary value ops/s.
- [ALR0507E](#) The Total Back-end I/O Rate of array, MDisk, or MDisk Group Array, MDisk, or MDisk Group name in storage system storage system name was measured to be measured value ops/s, which violated the defined critical-idle boundary value of boundary value ops/s.
- [ALR0508E](#) The Total Back-end Data Rate of array, MDisk, or MDisk Group Array, MDisk, or MDisk Group name in storage system storage system name was measured to be measured value MIB/s, which violated the defined critical-stress boundary value of boundary value MIB/s.
- [ALR0509W](#) The Total Back-end Data Rate of array, MDisk, or MDisk Group Array, MDisk, or MDisk Group name in storage system storage system name was measured to be measured value MIB/s, which violated the defined warning-stress boundary value of boundary value MIB/s.
- [ALR0510W](#) The Total Back-end Data Rate of array, MDisk, or MDisk Group Array, MDisk, or MDisk Group name in storage system storage system name was measured to be measured value MIB/s, which violated the defined warning-idle boundary value of boundary value MIB/s.
- [ALR0511E](#) The Total Back-end Data Rate of array, MDisk, or MDisk Group Array, MDisk, or MDisk Group name in storage system storage system name was measured to be measured value MIB/s, which violated the defined critical-idle boundary value of boundary value MIB/s.
- [ALR0512E](#) The Overall Back-end Response Time of MDisk MDisk name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-stress boundary value of boundary value ms/op.
- [ALR0513W](#) The Overall Back-end Response Time of MDisk MDisk name in storage system storage system name was measured to be measured value ms/op, which violated the defined warning-stress boundary value of boundary value ms/op.
- [ALR0514W](#) The Overall Back-end Response Time of MDisk MDisk name in storage system storage system name was measured to be measured value ms/op, which violated the defined warning-idle boundary value of boundary value ms/op.
- [ALR0515E](#) The Overall Back-end Response Time of MDisk MDisk name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-idle boundary value of boundary value ms/op.
- [ALR0516E](#) The Total I/O Rate of controller or I/O Group Controller or I/O Group name in storage system storage system name was measured to be measured value ops/s, which violated the defined critical-stress boundary value of boundary value ops/s.
- [ALR0517W](#) The Total I/O Rate of controller or I/O Group Controller or I/O Group name in storage system storage system name was measured to be measured value ops/s, which violated the defined warning-stress boundary value of boundary value ops/s.
- [ALR0518W](#) The Total I/O Rate of controller or I/O Group Controller or I/O Group name in storage system storage system name was measured to be measured value ops/s, which violated the defined warning-idle boundary value of boundary value ops/s.
- [ALR0519E](#) The Total I/O Rate of controller or I/O Group Controller or I/O Group name in storage system storage system name was measured to be measured value ops/s, which violated the defined critical-idle boundary value of boundary value ops/s.

- ## 1050 IBM Spectrum Control

- ## 1052 IBM Spectrum Control

- [ALR0712E The Zero Receive Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined critical-stress boundary value of boundary_value.](#)
- [ALR0713W The Zero Receive Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-stress boundary value of boundary_value.](#)
- [ALR0714W The Zero Receive Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined warning-idle boundary value of boundary_value.](#)
- [ALR0715E The Zero Receive Buffer Credit Percentage of port port_name in storage system device_name was found to be measured_value, which violated the defined critical-idle boundary value of boundary_value.](#)
- [ALR0716E The number of bytes received on each of the client network interfaces of cluster cluster_name in storage subsystem storage system name was measured to be measured_valuebytes, which violated the defined critical-idle boundary value of boundary_valuebytes.](#)
- [ALR1114I New Storage Subsystem discovered.](#)
- [ALR4000I Endpoint device endpoint device name has been discovered.](#)
- [ALR4001W Endpoint device endpoint device name is missing.](#)
- [ALR4002I Endpoint device endpoint device name has been rediscovered.](#)
- [ALR4015I subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN has been discovered.](#)
- [ALR4016W subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN is missing.](#)
- [ALR4017I subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN has been rediscovered.](#)
- [ALR4018W subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN has gone offline.](#)
- [ALR4019I subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN has gone online.](#)
- [ALR4020I Switch switch name or WWN has been discovered.](#)
- [ALR4021E Switch switch name or WWN is missing.](#)
- [ALR4022I Switch switch name or WWN has been rediscovered.](#)
- [ALR4023W The version for switch switch name or WWN has changed from previous firmware version to current firmware version.](#)
- [ALR4024W Status of switch switch name or WWN has degraded from previous status to current status.](#)
- [ALR4025I Status of switch switch name or WWN has improved from previous status to current status.](#)
- [ALR4026I Blade blade name or WWN on switch switch name or WWN has been discovered.](#)
- [ALR4027W Blade blade name or WWN on switch switch name or WWN is missing.](#)
- [ALR4028I Blade blade name or WWN on switch switch name or WWN has been rediscovered.](#)
- [ALR4029E Blade blade name or WWN on switch switch name or WWN has gone offline.](#)
- [ALR4030I Blade blade name or WWN on switch switch name or WWN has gone online.](#)
- [ALR4034W The driver version for HBA adapter name on server host name has changed from previous version to new version.](#)
- [ALR4035W The firmware version for HBA adapter name on server host name has changed from previous version to new version.](#)
- [ALR4046I Fabric fabric name or WWN is discovered.](#)
- [ALR4047E Fabric fabric name or WWN is missing.](#)
- [ALR4048I Fabric fabric name or WWN is rediscovered.](#)
- [ALR4051I Inactive zone zone name in fabric fabric name or WWN has been discovered.](#)
- [ALR4052W Inactive zone zone name in fabric fabric name or WWN is missing.](#)
- [ALR4053I Inactive zone zone name in fabric fabric name or WWN has been rediscovered.](#)
- [ALR4054I Inactive zoneset zoneset name in fabric fabric name or WWN has been discovered.](#)
- [ALR4055W Inactive zoneset zoneset name in fabric fabric name or WWN is missing.](#)
- [ALR4056I Inactive zoneset zoneset name in fabric fabric name or WWN has been rediscovered.](#)
- [ALR4063I The connection from switch or node switch name or WWN port port name or WWPN to switch or node switch name or WWN port port name or WWPN has been discovered.](#)
- [ALR4064W The connection from switch or node switch name or WWN port port name or WWPN to switch or node switch name or WWN port port name or WWPN is missing.](#)
- [ALR4065I The connection from switch or node switch name or WWN port port name or WWPN to switch or node switch name or WWN port port name or WWPN has been rediscovered.](#)
- [ALR4066I Switch switch name or WWN in fabric fabric name or WWN has been discovered.](#)
- [ALR4067W Switch switch name or WWN in fabric fabric name or WWN is missing.](#)
- [ALR4068I Switch switch name or WWN in fabric fabric name or WWN has been rediscovered.](#)
- [ALR4069I Port port name or WWPN in switch switch name or WWN has been discovered.](#)
- [ALR4070W Port port name or WWPN in switch switch name or WWN is missing.](#)
- [ALR4071I Port port name or WWPN in switch switch name or WWN has been rediscovered.](#)
- [ALR4078I Alias zone alias has been added to inactive zone zone name in fabric fabric name or WWN.](#)
- [ALR4079W Alias zone alias has been removed from inactive zone zone name in fabric fabric name or WWN.](#)
- [ALR4080I Alias zone alias has been readded to inactive zone zone name in fabric fabric name or WWN.](#)
- [ALR4081I Zone member zone member name has been added to inactive zone zone name in fabric fabric name or WWN.](#)
- [ALR4082I Zone member zone member name has been removed from inactive zone zone name in fabric fabric name or WWN.](#)
- [ALR4083I Zone member zone member name has been readded to inactive zone zone name in fabric fabric name or WWN.](#)
- [ALR4084I Zone zone name has been added to inactive zone set zone set name in fabric fabric name or WWN.](#)
- [ALR4085I Zone zone name has been removed from inactive zone set zone set name in fabric fabric name or WWN.](#)
- [ALR4086I Zone zone name has been readded to inactive zone set zone set name in fabric fabric name or WWN.](#)
- [ALR4089W ZoneSet zoneset name in fabric fabric name or WWN has been deactivated. ZoneSet zoneset name has been activated.](#)
- [ALR4090W Active zone zone name in fabric fabric name or WWN is missing.](#)
- [ALR4091I Active zone zone name in fabric fabric name or WWN has been discovered.](#)
- [ALR4092I Active zoneset zoneset name in fabric fabric name or WWN has been discovered.](#)
- [ALR4093E ZoneSet zoneset name in fabric fabric name or WWN has been deactivated. ZoneSet zoneset name has been activated.](#)
- [ALR4094I Active zone zone name in fabric fabric name or WWN has been rediscovered.](#)
- [ALR4095I Active zoneset zoneset name in fabric fabric name or WWN has been rediscovered.](#)
- [ALR4096I Zone member zone member name has been added to active zone zone name in fabric fabric name or WWN.](#)
- [ALR4097I Zone member zone member name has been removed from active zone zone name in fabric fabric name or WWN.](#)
- [ALR4098I Zone member zone member name has been readded to active zone zone name in fabric fabric name or WWN.](#)
- [ALR4099I Zone zone name has been added to active zone set zone set name in fabric fabric name or WWN.](#)
- [ALR4103W The performance monitor's primary process is experiencing a high memory utilization.](#)
- [ALR4104W A database used by the system is reporting an alarm value.](#)
- [ALR4105W Device server configuration should be changed to improve performance: value.](#)
- [ALR4106W The IBM Spectrum Control server is receiving a high number of external type of events received, which is either CIM for CIM indications or SNMP for SNMP traps events, which may cause temporary performance degradation.](#)
- [ALR4107I Zone zone name has been removed from active zone set zone set name in fabric fabric name or WWN.](#)
- [ALR4108I The server server name at host host name successfully connected to the database after previous attempts failed.](#)

- [ALR4112E The server server name at host host name failed to connect to the database.](#)
- [ALR4109I Alias zone alias has been added to active zone zone name in fabric fabric name or WWN.](#)
- [ALR4113E The IBM Spectrum Control server failed to connect to the database.](#)
- [ALR4110W Alias alias namehas been removed from active zone active zone name in fabric fabric name.](#)
- [ALR4111I Alias alias namehas been readded to active zone active zone name in fabric fabric name.](#)
- [ALR4197W A new connection is detected.](#)
- [ALR4198W The state for connection from initial state to final state has changed.](#)
- [ALR4199W The state for Fabric fabric name has changed.](#)
- [ALR4224W The state has changed for Node node name.](#)
- [ALR4225I Node node name has been discovered.](#)
- [ALR4226W Node node name has gone offline.](#)
- [ALR4227I Node node name has gone online.](#)
- [ALR4241E Subsystem storage subsystem name has gone offline.](#)
- [ALR4242I Subsystem storage subsystem name has gone online.](#)
- [ALR4243W The subsystem version has changed from initial version to new version on Subsystem storage subsystem name.](#)
- [ALR4244W The allocated capacity has changed from initial capacity to new capacity on Subsystem storage subsystem name.](#)
- [ALR4245W The available capacity has changed from initial capacity to new capacity on Subsystem storage subsystem name.](#)
- [ALR4246W Back-end capacity has changed from initial capacity to new capacity on Subsystem storage subsystem name.](#)
- [ALR4247W Back-end controller back-end controller name for owning storage subsystem name has gone offline.](#)
- [ALR4248I Back-end controller back-end controller name for owning storage subsystem name has gone online.](#)
- [ALR4249W Volume volume name on storage subsystem name has gone offline.](#)
- [ALR4250I Volume volume name on storage subsystem name has gone online.](#)
- [ALR4251W The capacity has changed from initial capacity to new capacity for Volume volume name on Subsystem storage subsystem name.](#)
- [ALR4252W The state for Pool pool name on Subsystem storage subsystem name has changed to not detectable.](#)
- [ALR4253I Pool pool name on Subsystem storage subsystem name has been discovered.](#)
- [ALR4254W Pool pool name on Subsystem storage subsystem name has gone offline.](#)
- [ALR4255I Pool pool name on Subsystem storage subsystem name has gone online.](#)
- [ALR4256W The pool capacity has changed from initial capacity to new capacity for Pool pool name on Subsystem storage subsystem name.](#)
- [ALR4257W The pool available space has changed from initial capacity to new capacity for Pool pool name on Subsystem storage subsystem name.](#)
- [ALR4273E Server server name has gone offline.](#)
- [ALR4274I Server server name has gone online.](#)
- [ALR4278W The property for Subsystem storage subsystem name has changed.](#)
- [ALR4300W The use count for Disk Drive disk drive nameon Subsystem subsystem namehas changed from initial use count to final use count.](#)
- [ALR4301W Disk Drive disk drive name on Subsystem subsystem namehas gone offline.](#)
- [ALR4304W The state for Back-end Controller back-end controller for subsystemhas changed from initial state to final state.](#)
- [ALR4305W The WWPN path count for Back-end Controller back-end controller for system name has changed from initial stateto final state.](#)
- [ALR4312W Notification has received from external device device name](#)
- [ALR4313W The endpoint version has changed from initial state to final stateon Endpoint](#)
- [ALR4314W Entity entity name has been discovered.](#)
- [ALR4315I Hypervisor hypervisor name has been discovered.](#)
- [ALR4316I Virtual machine vm name was added to hypervisor hypervisor name.](#)
- [ALR4317W Virtual machine vm name was removed from hypervisor hypervisor name.](#)
- [ALR4318W New unmanaged hypervisor discovered.](#)
- [ALR4319W Virtual machine added.](#)
- [ALR4320W Virtual machine removed.](#)
- [ALR4321W Hypervisor hypervisor missing.](#)
- [ALR4322W Hypervisor missing.](#)
- [ALR4323I New disk disk namediscovered for system system name.](#)
- [ALR4324W Disk disk not found for system system.](#)
- [ALR4325W New volume volume name discovered for system system name.](#)
- [ALR4326W Volume volume not found for system system.](#)
- [ALR4327W Zone Alias to Member Change](#)
- [ALR4328W The association between Zone Alias zone alias and Member member has changed.](#)
- [ALR4329I Zone member zone memberhas been added to zone alias zone aliasin fabric fabric.](#)
- [ALR4330W Zone member zone memberhas been removed from zone alias zone aliasin fabric fabric.](#)
- [ALR4331I Zone member zone memberhas been readded to zone alias zone aliasin fabric fabric.](#)
- [ALR4332W Storage Resource Agent Deployment Failed](#)
- [ALR4333W Replication Session State Change alert received.](#)
- [ALR4334W Replication Configuration Change alert received.](#)
- [ALR4335W Replication Suspending Event Notification alert received.](#)
- [ALR4336W Replication Communication Failure alert received.](#)
- [ALR4337W Replication Management Server State Change alert received](#)
- [ALR4338W Replication PPRC Path State Change alert received.](#)
- [ALR4339W The IBM Spectrum Control for Replication resource resource nametriggered an alert with the following message:message text](#)
- [ALR4353W number of affected datapaths Data Paths from Host host name to Volume volume name on Subsystem subsystem name are no longer available.](#)
- [ALR4354I number of affected datapaths Data Paths from Host host name to Volume volume name on Subsystem subsystem name have been discovered..](#)
- [ALR4356E The mount state of specified file system changed to error level.](#)
- [ALR4358I The mount state of specified file system changed to normal level.](#)
- [ALR4359E The CPU usage reached the error level.](#)
- [ALR4360W The CPU usage reached the warning level.](#)
- [ALR4361I The CPU usage reached the normal level.](#)
- [ALR4362E The memory usage reached the error level.](#)
- [ALR4364I The memory usage reached the normal level.](#)
- [ALR4365I The clustered CIFS is active.](#)
- [ALR4366W The clustered CIFS is disabled.](#)
- [ALR4367E The clustered CIFS status reached the error level.](#)
- [ALR4368I The IBM Spectrum Scale is active.](#)
- [ALR4369W The IBM Spectrum Scale status reached the warning level.](#)
- [ALR4370E The IBM Spectrum Scale is down.](#)
- [ALR4427I The file system file system has been detected on device type device device display name.](#)

- [ALR4429I The capacity of file system file system has changed from previous capacity to current capacity on device type device NAS display name.](#)
- [ALR4433W The free space on file system file system has fallen below the threshold value of threshold on device type device nas display name. The free space is current free space value or relative free space value of the file system capacity.](#)
- [ALR4440W The state of the device type node type node node name changed from old state to new state on device display name.](#)
- [ALR4441I The state of the device type node type node node name changed from old state to new state on device display name.](#)
- [ALR4442W The IBM Spectrum Scale status of device type node node name changed from old state to new state on device display name.](#)
- [ALR4443I The IBM Spectrum Scale status of device type node node name changed from old state to new state on device display name.](#)
- [ALR4447I The fileset fileset was detected on device type device device name.](#)
- [ALR4448I The fileset fileset was linked to path for device type device device name.](#)
- [ALR4455I The fileset fileset was unlinked on device type device device name.](#)
- [ALR4458W The number of free inodes on file system path has fallen below the threshold value of threshold on device type device device name. There are current value free inodes or current value relative to maximum of the maximum inodes.](#)
- [ALR4460I Export export name detected on device type device device name with path path.](#)
- [ALR4461W The state of export export name changed from previous state to current state on device type device device name.](#)
- [ALR4462W Export export name was reconfigured on device type device device name. Path changed from previous path to current path.](#)
- [ALR4463W Export export name was reconfigured on device type device device name. Protocols changed from previous list of protocols to current list of protocols.](#)
- [ALR4470W Export export name is missing from device type device device name..](#)
- [ALR4471I Export export name was rediscovered on device type device device name..](#)
- [ALR4474W Fileset fileset name is missing from device type device device name..](#)
- [ALR4475I Fileset fileset name was rediscovered on device type device device name..](#)
- [ALR4478W File system File system name is missing from device type device device name..](#)
- [ALR4479I File system File system name was rediscovered on device type device device name..](#)
- [ALR4482W A Quota type Quota limit type quota was violated for the path file system on the device device name device type system. Quota type resource name is consuming usage and the Quota limit type limit is threshold.](#)
- [ALR4496I New quota detected on file system path of Device type device Device name.](#)
- [ALR4385E The status of NSD NSD name reached error level.](#)
- [ALR4386W The status of NSD NSD name reached warning level.](#)
- [ALR4387I The status of NSD NSD name was set back to normal level.](#)
- [ALR4503I New NSD NSD name has been detected on device type device device display name.](#)
- [ALR4505W NSD NSD name is missing from device type device device name.](#)
- [ALR4507W The state of NSD NSD name changed from previous state to current state on device type device device name.](#)
- [ALR4511E Alert condition for nodes has been selected. Select only nodes.](#)
- [ALR4512E Alert condition for clusters has been selected. Select only clusters.](#)
- [ALR4513E Alert condition for NSD has been selected. Select only NSD.](#)
- [ALR4514E Alert condition for File set has been selected. Select only File set.](#)
- [ALR1022M A new unmanaged server or cluster discovered.](#)
- [ALR1294W The server or cluster has gone offline.](#)
- [ALR1295W The server or cluster has gone online.](#)
- [ALR1296W The server or cluster property has changed.](#)
- [ALR1245W A node state has changed.](#)
- [ALR1246W A node was discovered.](#)
- [ALR4528I Cluster was discovered.](#)
- [ALR4529I Cluster was removed.](#)
- [ALR4530I Cluster was rediscovered.](#)
- [ALR0078W =Performance monitor for device value failed to collect new data using data source value.](#)
- [ALR4391I Node node name is selected as cache gateway node.](#)
- [ALR4392I Node node name is unselected as cache gateway node.](#)
- [ALR4393I Home system home system name detected on device type device device name with path path.](#)
- [ALR4394I Home system has been removed from fileset fileset name on device type device device name.](#)
- [ALR4395W Home system home system name is missing from device type device device name.](#)
- [ALR4396I Home system home system name was rediscovered on device type device device name..](#)
- [ALR4397I Cache fileset cache fileset name detected on device type device device name.](#)
- [ALR4398W Cache fileset cache fileset name is missing from device type device device name.](#)
- [ALR4399I Cache fileset cache fileset name was rediscovered on device type device device name.](#)
- [ALR4400I Cache fileset name has changed from cache fileset name to cache fileset name on device type device device name.](#)
- [ALR4401I Cache fileset cache fileset name state has changed from old value to new value on device type device device name.](#)
- [ALR4402I Cache fileset cache fileset name mode has changed from old value to new value on device type device device name.](#)
- [ALR4403I Cache client cluster name is added to home system home system name on device type device device name.](#)
- [ALR4404I Cache client cluster name has been removed from home system home system name on device type device device name.](#)
- [ALR4541E The available space is too low for pool pool name on storage system storage system name. The measured value pool available space violates the critical boundary of user defined threshold value.](#)
- [ALR4542W The available space is too low for pool pool name on storage system storage system name. The measured value pool available space violates the warning boundary of user defined threshold value.](#)
- [ALR4543E The allocation is too high for pool pool name on storage system storage system name. The measured value pool virtual allocation violates the critical boundary of user defined threshold value.](#)
- [ALR4544W The allocation is too high for pool pool name on storage system storage system name. The measured value pool virtual allocation violates the warning boundary of user defined threshold value.](#)
- [ALR4545E The shortfall percentage is too high for pool pool name on storage system storage system name. The measured value pool shortfall percentage violates the critical boundary of user defined threshold value.](#)
- [ALR4546W The shortfall percentage is too high for pool pool name on storage system storage system name. The measured value pool shortfall percentage violates the warning boundary of user defined threshold value.](#)
- [ALR4547I VMWare Cluster cluster name discovered on hypervisor hypervisor name.](#)
- [ALR4548W VMWare Cluster cluster name removed from hypervisor hypervisor name.](#)
- [ALR4549I New cluster hypervisor relationship discovered.](#)
- [ALR4550W Cluster hypervisor relationship removed.](#)
- [ALR4551I Hypervisor hypervisor name was moved from VMWare Cluster old cluster nameto VMWare Cluster new cluster name.](#)
- [ALR4552I Cluster hypervisor relationship moved.](#)
- [ALR4600I Fabric Name fabric WWN changed to fabric WWN.](#)
- [ALR1349I A new path path name was discovered for disk disk name on host host name.](#)

- [ALR1350W The path path name was not found for disk disk name on host host name.](#)
- [ALR1351W The path path name for disk disk name on host host name is disconnected.](#)
- [ALR4604I The home system home system was linked to path on resource type resource resource name.](#)
- [ALR4605I The home system home system was unlinked from a path on resource type resource resource name.](#)
- [ALR1352E The status of disk disk name on server server name has degraded to Error from former status.](#)
- [ALR1353W The status of disk disk name on server server name has degraded to Warning from former status.](#)
- [ALR1354I The status of disk disk name on server server name has improved to Normal from former status.](#)
- [ALR4625I New entity type, entity name, added to system type system name.](#)

ALR0001I The amount of RAM on host *server name* has changed from *current value* to *new value*.

Explanation

The amount of RAM on the host server has changed from the current value to the new value.

ALR0002I The amount of virtual memory on host *server name* has changed from *current value* to *new value*.

Explanation

The amount of virtual memory on host server has changed from current value to the new value.

ALR0003I A new disk drive has been detected on host *server name*. Disk manufacturer/serial Number: *manufacturer/serial Number*.

Explanation

A new disk drive has been detected on the host server. The disk is identified by the specified Disk manufacturer/serial Number.

ALR0004E A previously visible disk drive can no longer be found on host *server name*. Disk manufacturer/serial number: *manufacturer/serial number*.

Explanation

A previously visible disk drive can no longer be found on the host server. The disk is identified by the specified Disk manufacturer/serial Number.

ALR0005I A new filesystem has been detected on host *server name*. Filesystem mount point: *mount point*.

Explanation

A new filesystem has been detected on the host server. The filesystem is identified by the specified Filesystem mount point.

ALR0006E A previously visible filesystem can no longer be found on host *server name*. Filesystem mount point: *mount point*.

Explanation

A previously visible filesystem can no longer be found on the host server. The filesystem is identified by the specified filesystem mount point.

ALR0007E A disk drive visible on host *server name* has predicted that a disk failure is imminent. Disk manufacturer/serial number: *manufacturer/serial number*.

Explanation

A disk drive visible on the host server has predicted that a disk failure is imminent. The disk is identified by the specified disk manufacturer/serial number.

ALR0008I The physical space definition of filesystem *filesystem name* on host *server name* has been reconfigured.

Explanation

The physical space definition of the specified filesystem on the specified host has been reconfigured.

ALR0009W The free space on filesystem *filesystem name* on host *server name* has fallen below the threshold value of *threshold*. The free space is *freespace* or *percent* of the filesystem capacity.

Explanation

The free space on identified filesystem on the host server has fallen below the threshold value that was designated.

ALR0010W The number of free inodes on filesystem *filesystem name* on host *server name* has fallen below the threshold value of *threshold*. The number of free inodes is *free inodes* or *percent* of the filesystem's total inodes.

Explanation

The number of free inodes on the identified filesystem on the identified host server has fallen below the threshold value that was designated.

ALR0011W A new grown defect has been detected on a disk visible to host *server name*. Disk manufacturer/serial Number: *disk manufacturer/serial number*, Current grown defects: *current*, Previous grown defects: *previous*.

Explanation

A new grown defect has been detected on a disk visible to the identified host server. The Disk is identified by the specified Disk manufacturer/serial Number.

ALR0012W The number of grown disk defects has exceeded the threshold value of *threshold*. Host: *server name*, Disk manufacturer/serial number: *disk manufacturer/serial number*, Current grown defects: *current*, Previous grown defects: *previous*.

Explanation

The number of grown disk defects has exceeded the threshold value defined.

ALR0013W A new monitored directory has been detected on host *server name*. Directory name: *directory name*, Directory Group: *directory group*.

Explanation

A new monitored directory has been detected on the identified host server.

ALR0014E A monitored directory has been removed from host *server name*. Directory name: *directory name*, Directory group: *directory group*.

Explanation

A monitored directory has been removed from the specified host server with the specified directory name and directory group.

ALR0015W Directory *directory* on host *server name* has exceeded its space usage quota of *quota*. The directory is currently consuming *usage* or *percent* of the filesystem capacity.

Explanation

The Directory specified on the host server has exceeded its space usage quota.

ALR0016W Filesystem *filesystem name* on host *server name* has violated a filesystem constraint. *number of files* file(s) consuming *space* or *percent* of the filesystem capacity are in violation of the conditions defined in this constraint. The constraint threshold is *threshold*. User *user name* has *number of files* files consuming *space* of storage. *violating owners*).

Explanation

The Filesystem specified on the host server has violated a filesystem constraint. A number of files that are consuming a specified amount or percentage of the filesystem capacity are in violation of the conditions defined in this constraint.

ALR0017E Host *server name* appears to be down. *number of attempts* attempt(s) to ping this host have failed.

Explanation

The Host server appears to be down. A number of attempts to ping this host have failed.

ALR0018W *quota name user* has exceeded a network storage usage quota of *usage*. This *user* is currently consuming *amount* of storage.

Explanation

A network storage usage quota has been exceeded.

ALR0019W *quota name user on host server name has exceeded a server storage usage quota of usage. This user is currently consuming amount of storage.*

Explanation

A server storage usage quota has been exceeded.

ALR0020W *quota name user on host server name has exceeded a filesystem usage quota of <usage>valueon filesystem filesystem name. This user is currently consuming amount of storage.*

Explanation

A filesystem usage quota has been exceeded.

ALR0021W *Run number number of job creator job name has failed on run number of total jobs total jobs.*

Explanation

A run of a job or schedule did not complete successfully on the specified servers. A run is a single invocation of a job. Depending on the type of job, this problem might occur if the IBM Spectrum Control servers are unavailable, the local area network is down, or communication with the target servers or Storage Resource agent cannot be established.

Action

Depending on the type of job, try the following actions to resolve the problem. Ensure that the IBM Spectrum Control servers are running and that the local area network is available. Verify that you have a network connection to the server on which the monitoring action will run. If a data collection job generated the warning, ensure that the Storage Resource agent on the target server is up and running. Check the log files of the servers for error messages that might help determine the problem. See the product information center for the location of these log files.

ALR0022I *Server server name has been discovered.*

Explanation

A new server has been discovered.

ALR0023W *Run number run number of job creator job name has failed.*

Explanation

A run of a job has failed.

ALR0024W *User user name on host server name has exceeded a table space usage quota of usage on table space rdbms type, table space instance. This user is currently consuming amount of storage.*

Explanation

The specified User on the host server has exceeded a table space usage quota on the specified table space.

ALR0025W User *user name* on host *server name* has exceeded an RDBMS instance usage quota of *quota* on *rdbms type* instance *instance name*. This user is currently consuming *amount* of storage.

Explanation

The specified User on the host server has exceeded an RDBMS instance usage quota.

ALR0026W User *user name* has exceeded a network database storage usage quota of *quota*. This user is currently consuming *amount* of storage.

Explanation

The specified User has exceeded a network database storage usage quota.

ALR0027W The *log directory file name* archived log directory value on host *server name* has exceeded the threshold value of *threshold*. This directory currently contains *number of logs* archived logs consuming *amount* of storage.

Explanation

The specified archived log directory on the host server has exceeded the specified threshold value.

ALR0028I A new *table space name* has been discovered on *rdbms type* instance on host *server name*. *rdbms instance: database*.

Explanation

A new table space has been discovered on the host server.

ALR0029E *value value* has been dropped. RDBMS: *value value*, Host: *host name*.

Explanation

The specified table space has been dropped.

ALR0030E *value value* has been taken offline. RDBMS: *value value*, Host: *host name*.

Explanation

The specified table space has been taken offline.

ALR0031W The free space on *value: value*, RDBMS: *value value*, host: *host name*, has fallen below the threshold value of *value*. The free space is *value* or *value* of the *value* capacity.

Explanation

The free space on the specified table space, RDBMS and host has fallen below the threshold value.

ALR0032W The free space on table space: *table space name*, RDBMS: *rdbms type instance*, database: *database*, host: *server name*, is fragmented across *number of extents* extents. This exceeds the threshold value of *threshold* extents. The largest contiguous free extent is *largest extent*.

Explanation

The free space on the specified table space, RDBMS, database and host is fragmented across multiple extents.

ALR0033W The largest free extent available on table space: *table space name*, RDBMS: *rdbms type rdbms name*, database: *database name*, host: *server name*, has fallen below the threshold value of *threshold*. The largest free extent is *largest extent*.

Explanation

The largest free extent available on table space, RDBMS, database and host has fallen below the threshold value.

ALR0034W Segment *segment name* of table/cluster *table name* on host: *server name*, RDBMS: *rdbms type rdbms name*, database: *database name*, is fragmented across *number of extents* extents. This exceeds the threshold value of *threshold* extents. This segment is a *value type* segment.

Explanation

A segment of the specified table/cluster on the specified host, RDBMS and database is fragmented across multiple extents.

ALR0035W Segment *value* of table/cluster *value* on host: *host name*, RDBMS: *value value*, database: *value*, is nearing the maximum number of extents available to it. This segment currently occupies *value* extents. The *value* additional extent(s) available to this segment falls below the defined threshold of *value* extent(s). The segment is a *type type* segment.

Explanation

The segment of the specified table/cluster on the specified host, RDBMS and database is nearing the maximum number of extents available to it.

ALR0036W Table/cluster *table name* on host: *server name*, RDBMS: *rdbms type rdbms name*, database: *database name*, has exceeded a space usage quota of *usage quota*. This table is currently consuming *amount* of storage.

Explanation

The Table/cluster on the specified host, RDBMS and database has exceeded a space usage quota.

ALR0037W Table/cluster *value* on host: *host name*, RDBMS: *value value*, database: *database name*, has exceeded a chained row quota of *value*. Statistics indicate that *value* rows or *value* of the total rows are chained.

Explanation

The Table/cluster on the specified host, RDBMS and database has exceeded a chained row quota.

ALR0038W Segment *name* of table/cluster *table name* on host: *server name*, RDBMS: *rd bms type rdbms name*, database: *database name*, has *amount* of unused, wasted space. This represents *amount* of the total space allocated to the segment, and exceeds the threshold value of *threshold*. This segment is a *type type* segment.

Explanation

The segment of the specified table/cluster on the specified host, RDBMS and database has wasted space.

ALR0039E Table/cluster *table name* on host: *server name*, RDBMS: *rd bms type rdbms name*, database: *database name*, has been dropped.

Explanation

The Table/cluster on the specified host, RDBMS and database has been dropped.

ALR0040W Filer *filer name* has been discovered. Spectrum Control will not monitor this filer until it has been licensed.

Explanation

The specified Filer has been discovered.

ALR0041W The amount of log freespace available on *instance: database*, RDBMS: *rd bms type* , *<rd bms>host: server name*, has fallen below the threshold value of *threshold*. The amount of log freespace available is *amount* or *percent* of the *total* capacity.

Explanation

The amount of log freespace available on the specified database, RDBMS and host has fallen below the threshold value.

ALR0042I A new device has been discovered on *rd bms* instance *instance* on host *server name*. Device: *device name*, Capacity: *capacity*, File Name: *filename*.

Explanation

A new device has been discovered on instance and host.

ALR0043E Device *device name* has been dropped from *rdbms* instance *instance* on host *server name*. Capacity: *capacity*, File Name: *filename*.

Explanation

The specified Device has been dropped from the specified instance and host.

ALR0044W The amount of device freespace available on *value* instance *value* on host *value* has fallen below the threshold value of *value*. The amount of device freespace available is *value* or *value* of the current capacity of *value*.

Explanation

The amount of device freespace available on the specified instance and host has fallen below the threshold value.

ALR0045W The amount of device freespace available on *device* instance *instance* on host *server name* has gone above the threshold value of *threshold*. The amount of device freespace available is *freespace* or *percent* of the current capacity of *total capacity*.

Explanation

The amount of device freespace available on the specified device on instance and host has gone above the threshold value.

ALR0046W Database *database name* has not been backed up in the last *number of days* days. Last backup for the database was on *date*.
RDBMS: *rdbms type* *rdbms name*, Host: *server name*.

Explanation

The specified Database has not been backed up in the last defined number of days.

ALR0047W *filer name* *filer type* is no longer accessible from *host entity*.

Explanation

The specified filer is no longer accessible.

ALR0048W Storage Subsystem *subsystem name* is no longer accessible from host *server name*.

Explanation

The specified storage subsystem is no longer accessible from the specified host server.

ALR0049W *disk array name* *disk array type* has been discovered from host *server name*. Spectrum Control will not monitor this *disk*

array until it has been selected as for monitoring from within the Storage Subsystem Administration GUI.

Explanation

The specified disk array has been discovered from the host server. Spectrum Control will not monitor this disk array until it has been selected for monitoring from within the storage subsystem Administration GUI.

ALR0050W The amount of cache on storage subsystem *subsystem name* has changed from *old value* to *new value*.

Explanation

The amount of Cache on the specified storage subsystem has changed.

ALR0052W Filesystem *filesystem name* on host *server name* will be automatically extended because its free space has fallen below the threshold of *threshold*. Current free space: *current free space*; Current capacity: *current capacity*; Target capacity: *target capacity* }.

Explanation

The specified Filesystem on the specified host will be automatically extended because its free space has fallen below the threshold defined.

ALR0053W Filesystem *filesystem name* on host *server name* needs extension but will not be because its current capacity of *current capacity* exceeds the specified limit of *limit*. Filesystem free space: *freespace (current capacity of current capacity)* .

Explanation

The specified Filesystem on the specified host needs to be extended but will not be because its current capacity exceeds the specified limit.

ALR0055I Cluster resource group *cluster resource group name* was added to cluster *cluster name* on node *node name* .

Explanation

The specified cluster resource group was added to the specified cluster on the specified node.

ALR0056I Cluster resource group *cluster resource group name* was removed from cluster *cluster name* on node *node name* .

Explanation

The specified cluster resource group was removed from the specified cluster on the specified node.

ALR0057I Cluster resource group *cluster resource group name* was moved in cluster *cluster name* from node *node name* to node *node name* .

Explanation

The specified cluster resource group was moved in the specified cluster to a different node.

ALR0076W Performance monitor failure for device value.

Explanation

The performance monitor for the specified device encountered an error and cannot continue collecting performance data.

Action

Check the performance monitor job log for additional information.

ALR0500E The Disk Utilization Percentage of array *array name* in storage system *storage system name* was measured to be *measured value*%. This violated the critical-stress boundary value of *boundary value*%.

Explanation

Disk Utilization Percentage measures the approximate utilization percentage of the arrays in monitored storage systems. This value represents the average percent of time that the disks associated with the array were busy. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array in the storage system has a Disk Utilization Percentage value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this array. Either the array is experiencing a hardware or software related problem that is causing a drop in performance, or the array is being overloaded with too high a workload. If the workload is too high, moving some of the array's workload to other less busy arrays might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the array. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

Note that some highly sequential workloads such as batch or backup processing might continually exceed the threshold because they drive the arrays to high utilization percentages. For these types of workloads, a high utilization indicates that the work is being performed very efficiently and is not a cause for concern. If this situation occurs for your workload, consider enabling the Sequential I/O Percentage filter. Use this filter to ignore any violations of the Disk Utilization Percentage threshold for highly-sequential workloads.

ALR0501W The Disk Utilization Percentage of array *array name* in storage system *storage system name* was measured to be *measured value*%. This violated the warning-stress boundary value of *boundary value*%.

Explanation

Disk Utilization Percentage measures the approximate utilization percentage of the arrays in monitored storage systems. This value represents the average percent of time that the disks associated with the array were busy. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array in the storage system has a Disk Utilization Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this array. Either the array is experiencing a hardware or software related problem that is causing a drop in performance, or the array is being overloaded with too high a workload. If the workload is too high, moving some of the array's workload to other less busy arrays might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the array. In this case, increase the boundary value in the threshold definition to reduce the

number of unnecessary alerts.

Note that some highly sequential workloads such as batch or backup processing might continually exceed the threshold because they drive the arrays to high utilization percentages. For these types of workloads, a high utilization indicates that the work is being performed very efficiently and is not a cause for concern. If this situation occurs for your workload, consider enabling the Sequential I/O Percentage filter. Use this filter to ignore any violations of the Disk Utilization Percentage threshold for highly-sequential workloads.

ALR0502W The Disk Utilization Percentage of array *array name* in storage system *storage system name* was measured to be *measured value%*. This violated the warning-idle boundary value of *boundary value%*.

Explanation

Disk Utilization Percentage measures the approximate utilization percentage of the arrays in monitored storage systems. This value represents the average percent of time that the disks associated with the array were busy. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array in the storage system has a Disk Utilization Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this array. This can happen if the application(s) using the array stopped working properly.

This type of threshold boundary should only be defined if the workload for the arrays in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0503E The Disk Utilization Percentage of array *array name* in storage system *storage system name* was measured to be *measured value%*. This violated the critical-idle boundary value of *boundary value%*.

Explanation

Disk Utilization Percentage measures the approximate utilization percentage of the arrays in monitored storage systems. This value represents the average percent of time that the disks associated with the array were busy. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array in the storage system has a Disk Utilization Percentage value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this array. This can happen if the application(s) using the array stopped working properly.

This type of threshold boundary should only be defined if the workload for the arrays in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0504E The Total Back-end I/O Rate of array, MDisk, or MDisk Group Array, *MDisk, or MDisk Group name* in storage system *storage system name* was measured to be *measured value ops/s*, which violated the defined critical-stress boundary value of *boundary value ops/s*.

Explanation

Total Back-end I/O Rate measures the average number of I/O operations per second for arrays, MDisks, and MDisk Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array, MDisk, or MDisk Group in the storage system has a Total Back-end I/O Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the array, MDisk, or MDisk Group is being overloaded with too high a workload. If this is the case, moving some of the array's or MDisk Group's workload to other less busy arrays or MDisk Groups might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the array, MDisk, or MDisk Group. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0505W The Total Back-end I/O Rate of array, MDisk, or MDisk Group *Array, MDisk, or MDisk Group name in storage system storage system name* was measured to be *measured value ops/s*, which violated the defined warning-stress boundary value of *boundary value ops/s*.

Explanation

Total Back-end I/O Rate measures the average number of I/O operations per second for arrays, MDisks, and MDisk Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array, MDisk, or MDisk Group in the storage system has a Total Back-end I/O Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the array, MDisk, or MDisk Group is being overloaded with too high a workload. If this is the case, moving some of the array's or MDisk Group's workload to other less busy arrays or MDisk Groups might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the array, MDisk, or MDisk Group. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0506W The Total Back-end I/O Rate of array, MDisk, or MDisk Group *Array, MDisk, or MDisk Group name in storage system storage system name* was measured to be *measured value ops/s*, which violated the defined warning-idle boundary value of *boundary value ops/s*.

Explanation

Total Back-end I/O Rate measures the average number of I/O operations per second for arrays, MDisks, and MDisk Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array, MDisk, or MDisk Group in the storage system has a Total Back-end I/O Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this array, MDisk, or MDisk Group. This can happen if the application(s) using the array, MDisk, or MDisk Group stopped working properly.

This type of threshold boundary should only be defined if the workload for the arrays, MDisks, or MDisk Groups in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0507E The Total Back-end I/O Rate of array, MDisk, or MDisk Group *Array, MDisk, or MDisk Group name in storage system storage system name* was measured to be *measured value ops/s*, which violated the defined critical-idle boundary value of *boundary value ops/s*.

Explanation

Total Back-end I/O Rate measures the average number of I/O operations per second for arrays, MDisks, and MDisk Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array, MDisk, or MDisk Group in the storage system has a Total Back-end I/O Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this array, MDisk, or MDisk Group. This can happen if the application(s) using the array, MDisk, or MDisk Group stopped working properly.

This type of threshold boundary should only be defined if the workload for the arrays, MDisks, or MDisk Groups in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0508E The Total Back-end Data Rate of array, MDisk, or MDisk Group *Array, MDisk, or MDisk Group name in storage system storage system name* was measured to be *measured value* MiB/s, which violated the defined critical-stress boundary value of *boundary value* MiB/s.

Explanation

Total Back-end Data Rate measures the average number of mebibytes per second transferred for arrays, MDisks, and MDisk Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array, MDisk, or MDisk Group in the storage system has a Total Back-end Data Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the array, MDisk, or MDisk Group is being overloaded with too high a workload. If this is the case, moving some of the array's or MDisk Group's workload to other less busy arrays or MDisk Groups might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the array, MDisk, or MDisk Group. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0509W The Total Back-end Data Rate of array, MDisk, or MDisk Group *Array, MDisk, or MDisk Group name in storage system storage system name* was measured to be *measured value* MiB/s, which violated the defined warning-stress boundary value of *boundary value* MiB/s.

Explanation

Total Back-end Data Rate measures the average number of mebibytes per second transferred for arrays, MDisks, and MDisk Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array, MDisk, or MDisk Group in the storage system has a Total Back-end Data Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the array, MDisk, or MDisk Group is being overloaded with too high a workload. If this is the case, moving some of the array's or MDisk Group's workload to other less busy arrays or MDisk Groups might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the array, MDisk, or MDisk Group. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0510W The Total Back-end Data Rate of array, MDisk, or MDisk Group Array, MDisk, or MDisk Group name in storage system storage system name was measured to be measured value MiB/s, which violated the defined warning-idle boundary value of boundary value MiB/s.

Explanation

Total Back-end Data Rate measures the average number of mebibytes per second transferred for arrays, MDisks, and MDisk Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array, MDisk, or MDisk Group in the storage system has a Total Back-end Data Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this array, MDisk, or MDisk Group. This can happen if the application(s) using the array, MDisk, or MDisk Group stopped working properly.

This type of threshold boundary should only be defined if the workload for the arrays, MDisks, or MDisk Groups in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0511E The Total Back-end Data Rate of array, MDisk, or MDisk Group Array, MDisk, or MDisk Group name in storage system storage system name was measured to be measured value MiB/s, which violated the defined critical-idle boundary value of boundary value MiB/s.

Explanation

Total Back-end Data Rate measures the average number of mebibytes per second transferred for arrays, MDisks, and MDisk Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array, MDisk, or MDisk Group in the storage system has a Total Back-end Data Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this array, MDisk, or MDisk Group. This can happen if the application(s) using the array, MDisk, or MDisk Group stopped working properly.

This type of threshold boundary should only be defined if the workload for the arrays, MDisks, or MDisk Groups in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0512E The Overall Back-end Response Time of MDisk MDisk name in storage system storage system name was measured to be measured value ms/op, which violated the defined critical-stress boundary value of boundary value ms/op.

Explanation

Overall Back-end Response Time measures the average number of milliseconds that it took to service each I/O operation for MDisks. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has an Overall Back-end Response Time that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this MDisk. Either the MDisk is experiencing a hardware or software related problem that is causing a drop in performance, or the array is being overloaded with too high a workload. If the workload is too high, moving some of the MDisk's workload to other less busy MDisk might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the MDisk. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

Due to internal cache management issues of some storage systems, it is possible that MDisk that are idle or almost idle have surprisingly large response times. Generally this is not a cause for concern, because application performance is rarely affected due to the low I/O rates. To avoid unnecessary threshold alerts for such cases, consider enabling the Total Back-end I/O Rate filter. Use this filter to ignore any violations of the Overall Back-end Response Time threshold for MDisk with very low I/O rates.

ALR0513W The Overall Back-end Response Time of MDisk *MDisk name* in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined warning-stress boundary value of *boundary value ms/op*.

Explanation

Overall Back-end Response Time measures the average number of milliseconds that it took to service each I/O operation for MDisk. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has an Overall Back-end Response Time that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this MDisk. Either the MDisk is experiencing a hardware or software related problem that is causing a drop in performance, or the array is being overloaded with too high a workload. If the workload is too high, moving some of the MDisk's workload to other less busy MDisk might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the MDisk. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

Due to internal cache management issues of some storage systems, it is possible that MDisk that are idle or almost idle have surprisingly large response times. Generally this is not a cause for concern, because application performance is rarely affected due to the low I/O rates. To avoid unnecessary threshold alerts for such cases, consider enabling the Total Back-end I/O Rate filter. Use this filter to ignore any violations of the Overall Back-end Response Time threshold for MDisk with very low I/O rates.

ALR0514W The Overall Back-end Response Time of MDisk *MDisk name* in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined warning-idle boundary value of *boundary value ms/op*.

Explanation

Overall Back-end Response Time measures the average number of milliseconds that it took to service each I/O operation for MDisk. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has an Overall Back-end Response Time that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this MDisk. This can happen if the application(s) using the MDisk stopped working properly.

This type of threshold boundary should only be defined if the workload for the MDisk in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0515E The Overall Back-end Response Time of MDisk *MDisk name* in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined critical-idle boundary value of *boundary value ms/op*.

Explanation

Overall Back-end Response Time measures the average number of milliseconds that it took to service each I/O operation for MDisk. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has an Overall Back-end Response Time that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this MDisk. This can happen if the application(s) using the MDisk stopped working properly.

This type of threshold boundary should only be defined if the workload for the MDisks in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0516E The Total I/O Rate of controller or I/O Group *Controller or I/O Group name* in storage system *storage system name* was measured to be *measured value ops/s*, which violated the defined critical-stress boundary value of *boundary value ops/s*.

Explanation

Total I/O Rate measures the average number of I/O operations per second for controllers and I/O Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or I/O Group in the storage system has a Total I/O Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the controller or I/O Group is being overloaded with too high a workload. If this is the case, moving some of the controller's or I/O Group's workload to other less busy controllers or I/O Groups might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the controller or I/O Group. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0517W The Total I/O Rate of controller or I/O Group *Controller or I/O Group name* in storage system *storage system name* was measured to be *measured value ops/s*, which violated the defined warning-stress boundary value of *boundary value ops/s*.

Explanation

Total I/O Rate measures the average number of I/O operations per second for controllers and I/O Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or I/O Group in the storage system has a Total I/O Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the controller or I/O Group is being overloaded with too high a workload. If this is the case, moving some of the controller's or I/O Group's workload to other less busy controllers or I/O Groups might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the controller or I/O Group. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0518W The Total I/O Rate of controller or I/O Group *Controller or I/O Group name* in storage system *storage system name* was measured to be *measured value ops/s*, which violated the defined warning-idle boundary value of *boundary value ops/s*.

Explanation

Total I/O Rate measures the average number of I/O operations per second for controllers and I/O Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or I/O Group in the storage system has a Total I/O Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this controller or I/O Group. This can happen if the application(s) using the controller or I/O Group stopped working properly.

This type of threshold boundary should only be defined if the workload for the controllers or I/O Groups in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0519E The Total I/O Rate of controller or I/O Group *Controller or I/O Group name* in storage system *storage system name* was measured to be *measured value ops/s*, which violated the defined critical-idle boundary value of *boundary value ops/s*.

Explanation

Total I/O Rate measures the average number of I/O operations per second for controllers and I/O Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or I/O Group in the storage system has a Total I/O Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this controller or I/O Group. This can happen if the application(s) using the controller or I/O Group stopped working properly.

This type of threshold boundary should only be defined if the workload for the controllers or I/O Groups in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0520E The Total Data Rate of controller or I/O Group *Controller or I/O Group name* in storage system *storage system name* was measured to be *measured value MiB/s*, which violated the defined critical-stress boundary value of *boundary value MiB/s*.

Explanation

Total Data Rate measures the average number of mebibytes per second transferred for controllers, and I/O Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or I/O Group in the storage system has a Total Data Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the controller or I/O Group is being overloaded with too high a workload. If this is the case, moving some of the controller's or I/O Group's workload to other less busy controllers or I/O Groups might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the controller or I/O Group. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0521W The Total Data Rate of controller or I/O Group *Controller or I/O Group name* in storage system *storage system name* was

measured to be *measured value* MiB/s, which violated the defined warning-stress boundary value of *boundary value* MiB/s.

Explanation

Total Data Rate measures the average number of mebibytes per second transferred for controllers and I/O Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or I/O Group in the storage system has a Total I/O Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the controller or I/O Group is being overloaded with too high a workload. If this is the case, moving some of the controller's or I/O Group's workload to other less busy controllers or I/O Groups might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the controller or I/O Group. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0522W The Total Data Rate of controller or I/O Group *Controller or I/O Group name* in storage system *storage system name* was measured to be *measured value* MiB/s, which violated the defined warning-idle boundary value of *boundary value* MiB/s.

Explanation

Total Data Rate measures the average number of mebibytes per second transferred for controllers and I/O Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or I/O Group in the storage system has a Total I/O Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this controller or I/O Group. This can happen if the application(s) using the controller or I/O Group stopped working properly.

This type of threshold boundary should only be defined if the workload for the controllers or I/O Groups in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0523E The Total Data Rate of controller or I/O Group *Controller or I/O Group name* in storage system *storage system name* was measured to be *measured value* MiB/s, which violated the defined critical-idle boundary value of *boundary value* MiB/s.

Explanation

Total Data Rate measures the average number of mebibytes per second transferred for controllers and I/O Groups. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or I/O Group in the storage system has a Total I/O Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this controller or I/O Group. This can happen if the application(s) using the controller or I/O Group stopped working properly.

This type of threshold boundary should only be defined if the workload for the controllers or I/O Groups in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0524E The Write-cache Delay Percentage of controller or node *controller or node name* in storage system *storage system name* was measured to be *measured value*%, which violated the critical-stress boundary value of *boundary value*%.

Explanation

Write-cache Delay Percentage measures the approximate percentage of I/O operations that were delayed due to write-cache space constraints or other conditions on the controllers and nodes in monitored storage systems. This value represents the average percent of all I/O operations that experienced delays. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or node in the storage system has a Write-cache Delay Percentage value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this controller or node. Either the amount of cache installed on the controller or node is insufficient to handle its workload, or the controller or node is experiencing a hardware or software related problem that is causing cache slots to be occupied longer than necessary due to a drop in write performance, or the controller or node is being overwhelmed with too high a write workload. If the workload is too high, moving some of the controller's or node's workload to other less busy controllers or nodes might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the controller or node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

Note there may be cases when the Write-cache Delay Percentage is high but the total number of delayed I/O operations is very low. Such cases are generally not a cause for concern, because application performance will rarely be affected due to the low I/O rates. To avoid unnecessary threshold alerts for such cases, consider enabling the Write-cache Delay I/O Rate filter. Use this filter to ignore any violations of the Write-cache Delay Percentage threshold for controllers or nodes with very low I/O rates.

ALR0525W The Write-cache Delay Percentage of controller or node *controller or node name* in storage system *storage system name* was measured to be *measured value*%, which violated the warning-stress boundary value of *boundary value*%.

Explanation

Write-cache Delay Percentage measures the approximate percentage of I/O operations that were delayed due to write-cache space constraints or other conditions on the controllers and nodes in monitored storage systems. This value represents the average percent of all I/O operations that experienced delays. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or node in the storage system has a Write-cache Delay Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this controller or node. Either the amount of cache installed on the controller or node is insufficient to handle its workload, or the controller or node is experiencing a hardware or software related problem that is causing cache slots to be occupied longer than necessary due to a drop in write performance, or the controller or node is being overwhelmed with too high a write workload. If the workload is too high, moving some of the controller's or node's workload to other less busy controllers or nodes might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the controller or node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

Note there may be cases when the Write-cache Delay Percentage is high but the total number of delayed I/O operations is very low. Such cases are generally not a cause for concern, because application performance will rarely be affected due to the low I/O rates. To avoid unnecessary threshold alerts for such cases, consider enabling the Write-cache Delay I/O Rate filter. Use this filter to ignore any violations of the Write-cache Delay Percentage threshold for controllers or nodes with very low I/O rates.

ALR0526W The Write-cache Delay Percentage of controller or node *controller or node name* in storage system *storage system name* was measured to be *measured value*%, which violated the warning-idle boundary value of *boundary value*%.

Explanation

Write-cache Delay Percentage measures the approximate percentage of I/O operations that were delayed due to write-cache space constraints or other conditions on the controllers and nodes in monitored storage systems. This value represents the average percent of all I/O operations that experienced delays. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or node in the storage system has a Write-cache Delay Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this controller or node. This can happen if the application(s) using the controller or node stopped working properly.

This type of threshold boundary should only be defined if the workload for the controllers or nodes in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0527E The Write-cache Delay Percentage of controller or node *controller or node name* in storage system *storage system name* was measured to be *measured value%*, which violated the critical-idle boundary value of *boundary value%*.

Explanation

Write-cache Delay Percentage measures the approximate percentage of I/O operations that were delayed due to write-cache space constraints or other conditions on the controllers and nodes in monitored storage systems. This value represents the average percent of all I/O operations that experienced delays. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller or node in the storage system has a Write-cache Delay Percentage value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this controller or node. This can happen if the application(s) using the controller or node stopped working properly.

This type of threshold boundary should only be defined if the workload for the controllers or nodes in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0528E The Cache Holding Time of controller *controller name* in storage system *storage system name* was measured to be *measured values*, which violated the critical-stress boundary value of *boundary values*.

Explanation

Cache Holding Time measures the average amount of time that data remains available in controller cache, in seconds. Shorter time periods indicate adverse performance. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller in the storage system has a Cache Holding Time value that is less than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this controller. Either the amount of cache installed on the controller is insufficient to handle its workload, or the controller is being overwhelmed with too high a workload. If the workload is too high, moving some of the controller's workload to other less busy controllers might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too high, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the controller. In this case, decrease the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0529W The Cache Holding Time of controller *controller name* in storage system *storage system name* was measured to be *measured*

values, which violated the warning-stress boundary value of boundary values.

Explanation

Cache Holding Time measures the average amount of time that data remains available in controller cache, in seconds. Shorter time periods indicate adverse performance. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller in the storage system has a Cache Holding Time value that is less than or equal to the warning-stress boundary, but greater than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this controller. Either the amount of cache installed on the controller is insufficient to handle its workload, or the controller is being overwhelmed with too high a workload. If the workload is too high, moving some of the controller's workload to other less busy controllers might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too high, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the controller. In this case, decrease the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0530W The Cache Holding Time of controller *controller name* in storage system *storage system name* was measured to be *measured values*, which violated the warning-idle boundary value of *boundary values*.

Explanation

Cache Holding Time measures the average amount of time that data remains available in controller cache, in seconds. Shorter time periods indicate adverse performance. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller in the storage system has a Cache Holding Time value that is greater than or equal to the warning-idle boundary, but less than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this controller. This can happen if the application(s) using the controller stopped working properly.

This type of threshold boundary should only be defined if the workload for the controllers in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider increasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0531E The Cache Holding Time of controller *controller name* in storage system *storage system name* was measured to be *measured values*, which violated the critical-idle boundary value of *boundary values*.

Explanation

Cache Holding Time measures the average amount of time that data remains available in controller cache, in seconds. Shorter time periods indicate adverse performance. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified controller in the storage system has a Cache Holding Time value that is greater than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this controller. This can happen if the application(s) using the controller stopped working properly.

This type of threshold boundary should only be defined if the workload for the controllers in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider

increasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0532E The Total Port I/O Rate of port *port name* in storage system *storage system name* was measured to be *measured value* ops/s, which violated the defined critical-stress boundary value of *boundary value* ops/s.

Explanation

Total Port I/O Rate measures the average number of I/O operations per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Total Port I/O Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the port is being overloaded with too high a workload. If this is the case, moving some of the port's workload to other less busy ports might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0533W The Total Port I/O Rate of port *port name* in storage system *storage system name* was measured to be *measured value* ops/s, which violated the defined warning-stress boundary value of *boundary value* ops/s.

Explanation

Total Port I/O Rate measures the average number of I/O operations per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Total Port I/O Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the port is being overloaded with too high a workload. If this is the case, moving some of the port's workload to other less busy ports might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0534W The Total Port I/O Rate of port *port name* in storage system *storage system name* was measured to be *measured value* ops/s, which violated the defined warning-idle boundary value of *boundary value* ops/s.

Explanation

Total Port I/O Rate measures the average number of I/O operations per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Total Port I/O Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0535E The Total Port I/O Rate of port *port name* in storage system *storage system name* was measured to be *measured value* ops/s, which violated the defined critical-idle boundary value of *boundary value* ops/s.

Explanation

Total Port I/O Rate measures the average number of I/O operations per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Total Port I/O Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0536E The Total Port Data Rate of port *port name* in device *device name* was measured to be *measured value* MiB/s, which violated the defined critical-stress boundary value of *boundary value* MiB/s.

Explanation

Total Port Data Rate measures the average number of mebibytes per second transferred for storage system ports or switch ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Total Port Data Rate that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the port is being overloaded with too high a workload. If this is the case, moving some of the port's workload to other less busy ports might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0537W The Total Port Data Rate of port *port name* in device *device name* was measured to be *measured value* MiB/s, which violated the defined warning-stress boundary value of *boundary value* MiB/s.

Explanation

Total Port Data Rate measures the average number of mebibytes per second transferred for storage system ports or switch ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Total Port Data Rate that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary, that was defined for the related threshold.

Action

The threshold violation might indicate that the port is being overloaded with too high a workload. If this is the case, moving some of the port's workload to other less busy ports might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0538W The Total Port Data Rate of port *port name* in device *device name* was measured to be *measured value* MiB/s, which violated the defined warning-idle boundary value of *boundary value* MiB/s.

Explanation

Total Port Data Rate measures the average number of mebibytes per second transferred for storage system ports or switch ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Total Port Data Rate that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary, that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0539E The Total Port Data Rate of port *port name* in device *device name* was measured to be *measured value* MiB/s, which violated the defined critical-idle boundary value of *boundary value* MiB/s.

Explanation

Total Port Data Rate measures the average number of mebibytes per second transferred for storage system ports or switch ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Total Port Data Rate that is less than or equal to the critical-idle boundary, that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0540E The Overall Port Response Time of port *port name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-stress boundary value of *boundary value* ms/op.

Explanation

Overall Port Response Time measures the average number of milliseconds that it took to service each I/O operation for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has an Overall Port Response Time that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this port. Either the port is experiencing a hardware or software related problem that is causing a drop in performance, or the port is being overloaded with too high a workload. If the workload is too high, moving some of the port's workload to other less busy ports might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0541W The Overall Port Response Time of port *port name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-stress boundary value of *boundary value* ms/op.

Explanation

Overall Port Response Time measures the average number of milliseconds that it took to service each I/O operation for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has an Overall Port Response Time that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this port. Either the port is experiencing a hardware or software related problem that is causing a drop in performance, or the port is being overloaded with too high a workload. If the workload is too high, moving some of the port's workload to other less busy ports might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0542W The Overall Port Response Time of port *port name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-idle boundary value of *boundary value* ms/op.

Explanation

Overall Port Response Time measures the average number of milliseconds that it took to service each I/O operation for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has an Overall Port Response Time that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0543E The Overall Port Response Time of port *port name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-idle boundary value of *boundary value* ms/op.

Explanation

Overall Port Response Time measures the average number of milliseconds that it took to service each I/O operation for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has an Overall Port Response Time that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0544E The Error Frame Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

Error Frame Rate measures the average number of frame errors that were received per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an Error Frame Rate that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the error rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0545W The Error Frame Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-stress boundary value of *boundary value* cnt/s.

Explanation

Error Frame Rate measures the average number of frame errors that were received per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an Error Frame Rate that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the error rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0546W The Error Frame Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

Error Frame Rate measures the average number of frame errors that were received per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an Error Frame Rate that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the port error thresholds, because a zero or low error rate is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0547E The Error Frame Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

Error Frame Rate measures the average number of frame errors that were received per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an Error Frame Rate that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the port error thresholds, because a zero or low error rate is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0548E The Link Failure Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

Link Failure Rate measures the average number of link errors that were received per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Failure Rate that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port, for the connected port, or for the fibre link or GBICs. Isolated errors can be ignored in most cases, but if the error rate remains consistently high over time, this may be a cause for concern. Ensure that both ports are enabled and configured properly, then follow the trouble-shooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0549W The Link Failure Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-stress boundary value of *boundary value* cnt/s.

Explanation

Link Failure Rate measures the average number of link errors that were received per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Failure Rate that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port, for the connected port, or for the fibre link or GBICs. Isolated errors can be ignored in most cases, but if the error rate remains consistently high over time, this may be a cause for concern. Ensure that both ports are enabled and configured properly, then follow the trouble-shooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0550W The Link Failure Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

Link Failure Rate measures the average number of link errors that were received per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Failure Rate that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the port error thresholds, because a zero or low error rate is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0551E The Link Failure Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

Link Failure Rate measures the average number of link errors that were received per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Failure Rate that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the port error thresholds, because a zero or low error rate is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0552E The Total Port Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* frm/s, which violated the defined critical-stress boundary value of *boundary value* frm/s.

Explanation

Total Port Frame Rate measures the average number of frames transferred per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Total Port Frame Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the port is being overloaded with too high a workload. If this is the case, moving some of the port's workload to other less busy ports might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0553W The Total Port Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* frm/s, which violated the defined warning-stress boundary value of *boundary value* frm/s.

Explanation

Total Port Frame Rate measures the average number of frames transferred per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Total Port Frame Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that the port is being overloaded with too high a workload. If this is the case, moving some of the port's workload to other less busy ports might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the workload for the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0554W The Total Port Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* frm/s, which violated the defined warning-idle boundary value of *boundary value* frm/s.

Explanation

Total Port Frame Rate measures the average number of frames transferred per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Total Port Frame Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0555E The Total Port Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* frm/s, which violated the defined critical-idle boundary value of *boundary value* frm/s.

Explanation

Total Port Frame Rate measures the average number of frames transferred per second for ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Total Port Frame Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0556E The CPU Utilization Percentage of node *node name* in storage system *storage system name* was measured to be *measured value*%, which violated the critical-stress boundary value of *boundary value*%.

Explanation

CPU Utilization Percentage measures the approximate utilization percentage of the CPUs of the nodes in the monitored storage systems. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a CPU Utilization Percentage value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this node. Either the node is experiencing a hardware or software related problem that is causing an increase in CPU usage, or the node is being overloaded with too high a workload. If the workload is too high, moving some of the node's workload to other less busy nodes might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0557W The CPU Utilization Percentage of node *node name* in storage system *storage system name* was measured to be *measured value%*, which violated the warning-stress boundary value of *boundary value%*.

Explanation

CPU Utilization Percentage measures the approximate utilization percentage of the CPUs of the nodes in the monitored storage systems. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a CPU Utilization Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this node. Either the node is experiencing a hardware or software related problem that is causing an increase in CPU usage, or the node is being overloaded with too high a workload. If the workload is too high, moving some of the node's workload to other less busy nodes might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0558W The CPU Utilization Percentage of node *node name* in storage system *storage system name* was measured to be *measured value%*, which violated the warning-idle boundary value of *boundary value%*.

Explanation

CPU Utilization Percentage measures the approximate utilization percentage of the CPUs of the nodes in the monitored storage systems. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a CPU Utilization Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the node stopped working properly.

This type of threshold boundary should only be defined if the workload for the nodes in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0559E The CPU Utilization Percentage of node *node name* in storage system *storage system name* was measured to be *measured*

value%, which violated the critical-idle boundary value of boundary value%.

Explanation

CPU Utilization Percentage measures the approximate utilization percentage of the CPUs of the nodes in the monitored storage systems. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a CPU Utilization Percentage value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the node stopped working properly.

This type of threshold boundary should only be defined if the workload for the nodes in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0560E The Back-end Read Response Time of array or MDisk array or MDisk name in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined critical-stress boundary value of *boundary value ms/op*.

Explanation

Back-end Read Response Time measures the average number of milliseconds that it took to service each read operation for arrays and MDisk. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array or MDisk in the storage system has a Back-end Read Response Time that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this array or MDisk. Either the array or MDisk is experiencing a hardware or software related problem that is causing a drop in performance, or the array or MDisk is being overloaded with too high a workload. If the workload is too high, moving some of the array or MDisk's workload to other less busy arrays or MDisk might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the MDisk. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

Due to internal cache management issues of some storage systems, it is possible that arrays or MDisk that are idle or almost idle have surprisingly large response times. Generally this is not a cause for concern, because application performance is rarely affected due to the low I/O rates. To avoid unnecessary threshold alerts for such cases, consider enabling the Back-end Read I/O Rate filter. Use this filter to ignore any violations of the Back-end Read Response Time threshold for arrays or MDisk with very low read I/O rates.

ALR0561W The Back-end Read Response Time of array or MDisk array or MDisk name in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined warning-stress boundary value of *boundary value ms/op*.

Explanation

Back-end Read Response Time measures the average number of milliseconds that it took to service each read operation for arrays and MDisk. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array or MDisk in the storage system has a Back-end Read Response Time that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this array or MDisk. Either the array or MDisk is experiencing a hardware or software related problem that is causing a drop in performance, or the array or MDisk is being overloaded with too high a workload. If the workload is too high, moving some of the array or MDisk's workload to other less busy arrays or MDisk might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the MDisk. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

Due to internal cache management issues of some storage systems, it is possible that arrays or MDisks that are idle or almost idle have surprisingly large response times. Generally this is not a cause for concern, because application performance is rarely affected due to the low I/O rates. To avoid unnecessary threshold alerts for such cases, consider enabling the Back-end Read I/O Rate filter. Use this filter to ignore any violations of the Back-end Read Response Time threshold for arrays or MDisks with very low read I/O rates.

ALR0562W The Back-end Read Response Time of array or MDisk array or MDisk name in storage system storage system name was measured to be *measured value ms/op*, which violated the defined warning-idle boundary value of *boundary value ms/op*.

Explanation

Back-end Read Response Time measures the average number of milliseconds that it took to service each read operation for arrays and MDisks. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array or MDisk in the storage system has a Back-end Read Response Time that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this array or MDisk. This can happen if the application(s) using the array or MDisk stopped working properly.

This type of threshold boundary should only be defined if the workload for the arrays or MDisks in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0563E The Back-end Read Response Time of array or MDisk array or MDisk name in storage system storage system name was measured to be *measured value ms/op*, which violated the defined critical-idle boundary value of *boundary value ms/op*.

Explanation

Back-end Read Response Time measures the average number of milliseconds that it took to service each read operation for arrays and MDisks. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array or MDisk in the storage system has a Back-end Read Response Time that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this array or MDisk. This can happen if the application(s) using the array or MDisk stopped working properly.

This type of threshold boundary should only be defined if the workload for the arrays or MDisks in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0564E The Back-end Write Response Time of array or MDisk array or MDisk name in storage system storage system name was measured to be *measured value ms/op*, which violated the defined critical-stress boundary value of *boundary value ms/op*.

Explanation

Back-end Write Response Time measures the average number of milliseconds that it took to service each write operation for arrays and MDisks. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array or MDisk in the storage system has a Back-end Write Response Time that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this array or MDisk. Either the array or MDisk is experiencing a hardware or software related problem that is causing a drop in performance, or the array or MDisk is being overloaded with too high a workload. If the workload is too high, moving some of the array or MDisk's workload to other less busy arrays or MDisks might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the MDisk. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

Due to internal cache management issues of some storage systems, it is possible that arrays or MDisks that are idle or almost idle have surprisingly large response times. Generally this is not a cause for concern, because application performance is rarely affected due to the low I/O rates. To avoid unnecessary threshold alerts for such cases, consider enabling the Back-end Write I/O Rate filter. Use this filter to ignore any violations of the Back-end Write Response Time threshold for arrays or MDisks with very low write I/O rates.

ALR0565W The Back-end Write Response Time of array or MDisk *array or MDisk name* in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined warning-stress boundary value of *boundary value ms/op*.

Explanation

Back-end Write Response Time measures the average number of milliseconds that it took to service each write operation for arrays and MDisks. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array or MDisk in the storage system has a Back-end Write Response Time that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this array or MDisk. Either the array or MDisk is experiencing a hardware or software related problem that is causing a drop in performance, or the array or MDisk is being overloaded with too high a workload. If the workload is too high, moving some of the array or MDisk's workload to other less busy arrays or MDisks might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the MDisk. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

Due to internal cache management issues of some storage systems, it is possible that arrays or MDisks that are idle or almost idle have surprisingly large response times. Generally this is not a cause for concern, because application performance is rarely affected due to the low I/O rates. To avoid unnecessary threshold alerts for such cases, consider enabling the Back-end Write I/O Rate filter. Use this filter to ignore any violations of the Back-end Write Response Time threshold for arrays or MDisks with very low write I/O rates.

ALR0566W The Back-end Write Response Time of array or MDisk *array or MDisk name* in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined warning-idle boundary value of *boundary value ms/op*.

Explanation

Back-end Write Response Time measures the average number of milliseconds that it took to service each write operation for arrays and MDisks. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array or MDisk in the storage system has a Back-end Write Response Time that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this array or MDisk. This can happen if the application(s) using the array or MDisk stopped working properly.

This type of threshold boundary should only be defined if the workload for the arrays or MDisks in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0567E The Back-end Write Response Time of array or MDisk *array or MDisk name* in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined critical-idle boundary value of *boundary value ms/op*.

Explanation

Back-end Write Response Time measures the average number of milliseconds that it took to service each write operation for arrays and MDisks. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified array or MDisk in the storage system has a Back-end Write Response Time that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this array or MDisk. This can happen if the application(s) using the array or MDisk stopped working properly.

This type of threshold boundary should only be defined if the workload for the arrays or MDisks in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0568E The Back-end Read Queue Time of MDisk *MDisk name* in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined critical-stress boundary value of *boundary value ms/op*.

Explanation

Back-end Read Queue Time measures the average number of milliseconds that each read operation for an MDisk spent on the queue before being issued to the back-end storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has a Back-end Read Queue Time that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this MDisk. Either the MDisk is experiencing a hardware or software related problem that is causing a drop in performance, or the MDisk is being overloaded with too high a workload. If the workload is too high, moving some of the MDisk's workload to other less busy MDisks might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the MDisk. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0569W The Back-end Read Queue Time of MDisk *MDisk name* in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined warning-stress boundary value of *boundary value ms/op*.

Explanation

Back-end Read Queue Time measures the average number of milliseconds that each read operation for an MDisk spent on the queue before being issued to the back-end storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has a Back-end Read Queue Time that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this MDisk. Either the MDisk is experiencing a hardware or software related problem that is causing a drop in performance, or the MDisk is being overloaded with too high a workload. If the workload is too high, moving some of the MDisk's workload to other less

busy MDisk might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the MDisk. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0570W The Back-end Read Queue Time of MDisk *MDisk name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-idle boundary value of *boundary value* ms/op.

Explanation

Back-end Read Queue Time measures the average number of milliseconds that each read operation for an MDisk spent on the queue before being issued to the back-end storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has a Back-end Read Queue Time that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this MDisk. This can happen if the application(s) using the array or MDisk stopped working properly.

It makes little sense to enable an idle boundary for the queue time thresholds, because a zero or low queue time is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0571E The Back-end Read Queue Time of MDisk *MDisk name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-idle boundary value of *boundary value* ms/op.

Explanation

Back-end Read Queue Time measures the average number of milliseconds that each read operation for an MDisk spent on the queue before being issued to the back-end storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has a Back-end Read Queue Time that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this MDisk. This can happen if the application(s) using the array or MDisk stopped working properly.

It makes little sense to enable an idle boundary for the queue time thresholds, because a zero or low queue time is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0572E The Back-end Write Queue Time of MDisk *MDisk name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-stress boundary value of *boundary value* ms/op.

Explanation

Back-end Write Queue Time measures the average number of milliseconds that each write operation for an MDisk spent on the queue before being issued to the back-end storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has a Back-end Write Queue Time that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this MDisk. Either the MDisk is experiencing a hardware or software related problem that is causing a drop in performance, or the MDisk is being overloaded with too high a workload. If the workload is too high, moving some of the MDisk's workload to other less busy MDisk might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the MDisk. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0573W The Back-end Write Queue Time of MDisk *MDisk name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-stress boundary value of *boundary value* ms/op.

Explanation

Back-end Write Queue Time measures the average number of milliseconds that each write operation for an MDisk spent on the queue before being issued to the back-end storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has a Back-end Write Queue Time that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this MDisk. Either the MDisk is experiencing a hardware or software related problem that is causing a drop in performance, or the MDisk is being overloaded with too high a workload. If the workload is too high, moving some of the MDisk's workload to other less busy MDisk might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the MDisk. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0574W The Back-end Write Queue Time of MDisk *MDisk name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-idle boundary value of *boundary value* ms/op.

Explanation

Back-end Write Queue Time measures the average number of milliseconds that each write operation for an MDisk spent on the queue before being issued to the back-end storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has a Back-end Write Queue Time that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this MDisk. This can happen if the application(s) using the array or MDisk stopped working properly.

It makes little sense to enable an idle boundary for the queue time thresholds, because a zero or low queue time is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0575E The Back-end Write Queue Time of MDisk *MDisk name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-idle boundary value of *boundary value* ms/op.

Explanation

Back-end Write Queue Time measures the average number of milliseconds that each write operation for an MDisk spent on the queue before being issued to the back-end storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified MDisk in the storage system has a Back-end Write Queue Time that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this MDisk. This can happen if the application(s) using the array or MDisk stopped working properly.

It makes little sense to enable an idle boundary for the queue time thresholds, because a zero or low queue time is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0576E The Port to Local Node Send Response Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-stress boundary value of *boundary value* ms/op. Violation of this threshold boundary could mean that it is taking too long to send data between nodes on the fabric, which suggests either a problem with the nodes or congestion around the associated FC ports on the fabric.

Explanation

Port to Local Node Send Response Time measures the average number of milliseconds that it took to service each send operation to another node in the local storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Send Response Time that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this node. Either the node is experiencing a hardware or software related problem that is causing a drop in performance, or there is some congestion around the storage system's ports in the fabric, which is causing unreasonable delays. Check for errors on the ports, and if necessary follow your storage system's troubleshooting guidelines to determine the root of the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0577W The Port to Local Node Send Response Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-stress boundary value of *boundary value* ms/op. Violation of this threshold boundary could mean that it is taking too long to send data between nodes on the fabric, which suggests either a problem with the nodes or congestion around the associated FC ports on the fabric.

Explanation

Port to Local Node Send Response Time measures the average number of milliseconds that it took to service each send operation to another node in the local storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Send Response Time that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this node. Either the node is experiencing a hardware or software related problem that is causing a drop in performance, or there is some congestion around the storage system's ports in the fabric, which is causing unreasonable delays. Check for errors on the ports, and if necessary follow your storage system's troubleshooting guidelines to determine the root of the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0578W The Port to Local Node Send Response Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-idle boundary value of *boundary value* ms/op.

Explanation

Port to Local Node Send Response Time measures the average number of milliseconds that it took to service each send operation to another node in the local storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Send Response Time that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the node stopped working properly.

This type of threshold boundary should only be defined if the workload for the nodes in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0579E The Port to Local Node Send Response Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-idle boundary value of *boundary value* ms/op.

Explanation

Port to Local Node Send Response Time measures the average number of milliseconds that it took to service each send operation to another node in the local storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Send Response Time that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the node stopped working properly.

This type of threshold boundary should only be defined if the workload for the nodes in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0580E The Port to Local Node Receive Response Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-stress boundary value of *boundary value* ms/op. Violation of this threshold boundary could mean that it is taking too long to send data between nodes on the fabric, which suggests either a problem with the nodes or congestion around the associated FC ports on the fabric.

Explanation

Port to Local Node Receive Response Time measures the average number of milliseconds that it took to service each receive operation from another node in the local storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Receive Response Time that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this node. Either the node is experiencing a hardware or software related problem that is causing a drop in performance, or there is some congestion around the storage system's ports in the fabric, which is causing unreasonable delays. Check for errors on the ports, and if necessary follow your storage system's troubleshooting guidelines to determine the root of the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0581W The Port to Local Node Receive Response Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-stress boundary value of *boundary value* ms/op. Violation of this threshold boundary could mean that it is taking too long to send data between nodes on the fabric, which suggests either a problem with the nodes or congestion around the associated FC ports on the fabric.

Explanation

Port to Local Node Receive Response Time measures the average number of milliseconds that it took to service each receive operation from another node in the local storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Receive Response Time that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this node. Either the node is experiencing a hardware or software related problem that is causing a drop in performance, or there is some congestion around the storage system's ports in the fabric, which is causing unreasonable delays. Check for errors on the ports, and if necessary follow your storage system's troubleshooting guidelines to determine the root of the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0582W The Port to Local Node Receive Response Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-idle boundary value of *boundary value* ms/op.

Explanation

Port to Local Node Receive Response Time measures the average number of milliseconds that it took to service each receive operation from another node in the local storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Receive Response Time that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the node stopped working properly.

This type of threshold boundary should only be defined if the workload for the nodes in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0583E The Port to Local Node Receive Response Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-idle boundary value of *boundary value* ms/op.

Explanation

Port to Local Node Receive Response Time measures the average number of milliseconds that it took to service each receive operation from another node in the local storage system. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Receive Response Time that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the node stopped working properly.

This type of threshold boundary should only be defined if the workload for the nodes in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0584E The Port to Local Node Send Queue Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-stress boundary value of *boundary value* ms/op. Violation of this threshold boundary could mean that the node has to wait too long to send data to other nodes on the fabric, which suggests congestion around the associated FC ports on the fabric.

Explanation

Port to Local Node Send Queue Time measures the average number of milliseconds that each send operation to another node in the local storage system spent on the queue before being issued. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Send Queue Time that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this node. Either the node is experiencing a hardware or software related problem that is causing a drop in performance, or there is some congestion around the storage system's ports in the fabric, which is causing unreasonable delays. Check for errors on the ports, and if necessary follow your storage system's troubleshooting guidelines to determine the root of the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0585W The Port to Local Node Send Queue Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-stress boundary value of *boundary value* ms/op. Violation of this threshold boundary could mean that the node has to wait too long to send data to other nodes on the fabric, which suggests congestion around the associated FC ports on the fabric.

Explanation

Port to Local Node Send Queue Time measures the average number of milliseconds that each send operation to another node in the local storage system spent on the queue before being issued. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Send Queue Time that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this node. Either the node is experiencing a hardware or software related problem that is causing a drop in performance, or there is some congestion around the storage system's ports in the fabric, which is causing unreasonable delays. Check for errors on the ports, and if necessary follow your storage system's troubleshooting guidelines to determine the root of the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0586W The Port to Local Node Send Queue Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-idle boundary value of *boundary value* ms/op.

Explanation

Port to Local Node Send Queue Time measures the average number of milliseconds that each send operation to another node in the local storage system spent on the queue before being issued. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Send Queue Time that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the node stopped working properly.

It makes little sense to enable an idle boundary for the queue time thresholds, because a zero or low queue time is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0587E The Port to Local Node Send Queue Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-idle boundary value of *boundary value* ms/op.

Explanation

Port to Local Node Send Queue Time measures the average number of milliseconds that each send operation to another node in the local storage system spent on the queue before being issued. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Send Queue Time that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the node stopped working properly.

It makes little sense to enable an idle boundary for the queue time thresholds, because a zero or low queue time is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0588E The Port to Local Node Receive Queue Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-stress

boundary value of *boundary value ms/op*. Violation of this threshold boundary could mean that the node has to wait too long to receive data to other nodes on the fabric, which suggests congestion around the associated FC ports on the fabric.

Explanation

Port to Local Node Receive Queue Time measures the average number of milliseconds that each receive operation to another node in the local storage system spent on the queue before being issued. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Receive Queue Time that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this node. Either the node is experiencing a hardware or software related problem that is causing a drop in performance, or there is some congestion around the storage system's ports in the fabric, which is causing unreasonable delays. Check for errors on the ports, and if necessary follow your storage system's troubleshooting guidelines to determine the root of the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0589W The Port to Local Node Receive Queue Time of node *node name* in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined warning-stress boundary value of *boundary value ms/op*. Violation of this threshold boundary could mean that the node has to wait too long to receive data to other nodes on the fabric, which suggests congestion around the associated FC ports on the fabric.

Explanation

Port to Local Node Receive Queue Time measures the average number of milliseconds that each receive operation to another node in the local storage system spent on the queue before being issued. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Receive Queue Time that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this node. Either the node is experiencing a hardware or software related problem that is causing a drop in performance, or there is some congestion around the storage system's ports in the fabric, which is causing unreasonable delays. Check for errors on the ports, and if necessary follow your storage system's troubleshooting guidelines to determine the root of the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0590W The Port to Local Node Receive Queue Time of node *node name* in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined warning-idle boundary value of *boundary value ms/op*.

Explanation

Port to Local Node Receive Queue Time measures the average number of milliseconds that each receive operation to another node in the local storage system spent on the queue before being issued. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Receive Queue Time that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the node stopped working properly.

It makes little sense to enable an idle boundary for the queue time thresholds, because a zero or low queue time is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0591E The Port to Local Node Receive Queue Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-idle boundary value of *boundary value* ms/op.

Explanation

Port to Local Node Receive Queue Time measures the average number of milliseconds that each receive operation to another node in the local storage system spent on the queue before being issued. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Port to Local Node Receive Queue Time that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the node stopped working properly.

It makes little sense to enable an idle boundary for the queue time thresholds, because a zero or low queue time is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0592E The Non-preferred Node Usage Percentage of I/O Group *I/O Group name* in device *device name* was measured to be *measured value*%, which violated the critical-stress boundary value of *boundary value*%.

Explanation

Non-preferred Node Usage Percentage measures the approximate percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified I/O Group in the storage system has a Non-preferred Node Usage Percentage value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

Under normal conditions, the non-preferred node usage percentage should be zero for all I/O groups. The threshold violation might indicate either that a configuration problem exists for some of the host servers using storage from the storage system, or that some of the data paths from the host servers to the storage system through the fabric are currently not operational.

To determine which volumes are the culprit, the Non-preferred Node Usage Percentage metric is also available for volumes. The hosts that are using these volumes may not have multi-pathing software installed or configured properly, such that volume I/O is being directed to the non-preferred node of the I/O Group. Or if the host does not have redundant data paths to the storage system, it is possible that the host was connected or zoned to the incorrect node of the I/O Group. Correcting these configuration errors will usually resolve the problem.

ALR0593W The Non-preferred Node Usage Percentage of I/O Group *I/O Group name* in device *device name* was measured to be *measured value*%, which violated the warning-stress boundary value of *boundary value*%.

Explanation

Non-preferred Node Usage Percentage measures the approximate percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified I/O Group in the storage system has a Non-preferred Node Usage Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

Under normal conditions, the non-preferred node usage percentage should be zero for all I/O groups. The threshold violation might indicate either that a configuration problem exists for some of the host servers using storage from the storage system, or that some of the data paths from the host servers to the storage system through the fabric are currently not operational.

To determine which volumes are the culprit, the Non-preferred Node Usage Percentage metric is also available for volumes. The hosts that are using these volumes may not have multi-pathing software installed or configured properly, such that volume I/O is being directed to the non-preferred node of the I/O Group. Or if the host does not have redundant data paths to the storage system, it is possible that the host was connected or zoned to the incorrect node of the I/O Group. Correcting these configuration errors will usually resolve the problem.

ALR0594W The Non-preferred Node Usage Percentage of I/O Group *I/O Group name* in device *device name* was measured to be *measured value%*, which violated the warning-idle boundary value of *boundary value%*.

Explanation

Non-preferred Node Usage Percentage measures the approximate percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified I/O Group in the storage system has a Non-preferred Node Usage Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for this type of threshold, because a zero or low non-preferred node usage percentage is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0595E The Non-preferred Node Usage Percentage of I/O Group *I/O Group name* in device *device name* was measured to be *measured value%*, which violated the critical-idle boundary value of *boundary value%*.

Explanation

Non-preferred Node Usage Percentage measures the approximate percentage of I/O operations that are not directed against the preferred node for each volume in an I/O Group. There is a small performance penalty when I/O does not go to the preferred node for each volume. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified I/O Group in the storage system has a Non-preferred Node Usage Percentage value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for this type of threshold, because a zero or low non-preferred node usage percentage is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0596E The Peak Back-end Write Response Time of node *node name* in storage system *storage system name* was measured to be *measured value ms/op*, which violated the defined critical-stress boundary value of *boundary value ms/op*. If writes to disk are too slow, and writes are being received faster than they can be destaged to disk, then the nodes write-cache will eventually fill up. In

extreme cases, the node will stop caching write data, causing a significant performance degradation for the affected volumes.

Explanation

Peak Back-end Write Response Time measures the peak (worst) number of milliseconds that it took to service each MDisk write operation by a node. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Peak Back-end Write Response Time that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for one or more MDisks accessed by this node. To determine which of the MDisks are the culprits, this metric is also available for MDisks. Especially when the peak back-end write response time is in the dozens of seconds, it is likely that write requests are being received faster than they can be destaged to disk. Because the cache on the nodes is not infinite, eventually the storage system will have to switch to a mode where write I/Os are no longer cached for the MDisk, which can have a significant impact on application performance.

Possible causes for this condition might be that some of the MDisks are experiencing hardware or software related problems that are causing a drop in performance, or that the MDisk Group is being overloaded with too high a workload. If the workload is too high, moving some of the MDisk Group's workload to other less busy MDisk Groups might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0597W The Peak Back-end Write Response Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-stress boundary value of *boundary value* ms/op. If writes to disk are too slow, and writes are being received faster than they can be destaged to disk, then the nodes write-cache will eventually fill up. In extreme cases, the node will stop caching write data, causing a significant performance degradation for the affected volumes.

Explanation

Peak Back-end Write Response Time measures the peak (worst) number of milliseconds that it took to service each MDisk write operation by a node. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Peak Back-end Write Response Time that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for one or more MDisks accessed by this node. To determine which of the MDisks are the culprits, this metric is also available for MDisks. Especially when the peak back-end write response time is in the dozens of seconds, it is likely that write requests are being received faster than they can be destaged to disk. Because the cache on the nodes is not infinite, eventually the storage system will have to switch to a mode where write I/Os are no longer cached for the MDisk, which can have a significant impact on application performance.

Possible causes for this condition might be that some of the MDisks are experiencing hardware or software related problems that are causing a drop in performance, or that the MDisk Group is being overloaded with too high a workload. If the workload is too high, moving some of the MDisk Group's workload to other less busy MDisk Groups might resolve the problem.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the node. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0598W The Peak Back-end Write Response Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined warning-idle boundary value of *boundary value* ms/op.

Explanation

Peak Back-end Write Response Time measures the peak (worst) number of milliseconds that it took to service each MDisk write operation by a node. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Peak Back-end Write Response Time that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the array or MDisk stopped working properly.

This type of threshold boundary should only be defined if the workload for the nodes in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0599E The Peak Back-end Write Response Time of node *node name* in storage system *storage system name* was measured to be *measured value* ms/op, which violated the defined critical-idle boundary value of *boundary value* ms/op.

Explanation

Peak Back-end Write Response Time measures the peak (worst) number of milliseconds that it took to service each MDisk write operation by a node. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified node in the storage system has a Peak Back-end Write Response Time that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this node. This can happen if the application(s) using the array or MDisk stopped working properly.

This type of threshold boundary should only be defined if the workload for the nodes in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0600E The Port Send Utilization Percentage of port *port name* in storage system *storage system name* was measured to be *measured value*%, which violated the critical-stress boundary value of *boundary value*%.

Explanation

Port Send Utilization Percentage measures the approximate utilization percentage of the ports in monitored storage systems. This value represents the average percent of time that the ports were busy sending data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Port Send Utilization Percentage value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this port. Either the port is experiencing a hardware or software related problem that is causing a drop in performance, or the port is being overloaded with too high a workload. If the workload is too high, moving some of the port's workload to other less busy ports might resolve the problem.

If the Port Send Bandwidth Percentage metric is low or within normal bounds, but the Port Send Utilization Percentage metric is high for this port, then there might be some downstream congestion in the fabric along the datapath for this port. Check to make sure that the related switch ports are operational, and that the server HBA and the related switch ports have similar hardware capabilities as the subsystem port. Less capable ports at the server HBA or at the switches along the data path may artificially restrict the speed, so communication can only go as fast as the slowest link on the path.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0601W The Port Send Utilization Percentage of port *port name* in storage system *storage system name* was measured to be *measured value%*, which violated the warning-stress boundary value of *boundary value%*.

Explanation

Port Send Utilization Percentage measures the approximate utilization percentage of the ports in monitored storage systems. This value represents the average percent of time that the ports were busy sending data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Port Send Utilization Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this port. Either the port is experiencing a hardware or software related problem that is causing a drop in performance, or the port is being overloaded with too high a workload. If the workload is too high, moving some of the port's workload to other less busy ports might resolve the problem.

If the Port Send Bandwidth Percentage metric is low or within normal bounds, but the Port Send Utilization Percentage metric is high for this port, then there might be some downstream congestion in the fabric along the datapath for this port. Check to make sure that the related switch ports are operational, and that the server HBA and the related switch ports have similar hardware capabilities as the subsystem port. Less capable ports at the server HBA or at the switches along the data path may artificially restrict the speed, so communication can only go as fast as the slowest link on the path.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0602W The Port Send Utilization Percentage of port *port name* in storage system *storage system name* was measured to be *measured value%*, which violated the warning-idle boundary value of *boundary value%*.

Explanation

Port Send Utilization Percentage measures the approximate utilization percentage of the ports in monitored storage systems. This value represents the average percent of time that the ports were busy sending data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Port Send Utilization Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0603E The Port Send Utilization Percentage of port *port name* in storage system *storage system name* was measured to be *measured value%*, which violated the critical-idle boundary value of *boundary value%*.

Explanation

Port Send Utilization Percentage measures the approximate utilization percentage of the ports in monitored storage systems. This value represents the average percent of time that the ports were busy sending data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Port Send Utilization Percentage value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0604E The Port Receive Utilization Percentage of port *port name* in storage system *storage system name* was measured to be *measured value%*, which violated the critical-stress boundary value of *boundary value%*.

Explanation

Port Receive Utilization Percentage measures the approximate utilization percentage of the ports in monitored storage systems. This value represents the average percent of time that the ports were busy receiving data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Port Receive Utilization Percentage value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this port. Either the port is experiencing a hardware or software related problem that is causing a drop in performance, or the port is being overloaded with too high a workload. If the workload is too high, moving some of the port's workload to other less busy ports might resolve the problem.

If the Port Receive Bandwidth Percentage metric is low or within normal bounds, but the Port Receive Utilization Percentage metric is high for this port, then there might be some downstream congestion in the fabric along the datapath for this port. Check to make sure that the related switch ports are operational, and that the server HBA and the related switch ports have similar hardware capabilities as the subsystem port. Less capable ports at the server HBA or at the switches along the data path may artificially restrict the speed, so communication can only go as fast as the slowest link on the path.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0605W The Port Receive Utilization Percentage of port *port name* in storage system *storage system name* was measured to be *measured value%*, which violated the warning-stress boundary value of *boundary value%*.

Explanation

Port Receive Utilization Percentage measures the approximate utilization percentage of the ports in monitored storage systems. This value represents the average percent of time that the ports were busy receiving data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Port Receive Utilization Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this port. Either the port is experiencing a hardware or software related problem that is causing a drop in performance, or the port is being overloaded with too high a workload. If the workload is too high, moving some of the port's workload to other less busy ports might resolve the problem.

If the Port Receive Bandwidth Percentage metric is low or within normal bounds, but the Port Receive Utilization Percentage metric is high for this port, then there might be some downstream congestion in the fabric along the datapath for this port. Check to make sure that the related switch ports are operational, and that the server HBA and the related switch ports have similar hardware capabilities as the subsystem port. Less capable ports at the server HBA or at the switches along the data path may artificially restrict the speed, so communication can only go as fast as the slowest link on the path.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0606W The Port Receive Utilization Percentage of port *port name* in storage system *storage system name* was measured to be *measured value%*, which violated the warning-idle boundary value of *boundary value%*.

Explanation

Port Receive Utilization Percentage measures the approximate utilization percentage of the ports in monitored storage systems. This value represents the average percent of time that the ports were busy receiving data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Port Receive Utilization Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0607E The Port Receive Utilization Percentage of port *port name* in storage system *storage system name* was measured to be *measured value%*, which violated the critical-idle boundary value of *boundary value%*.

Explanation

Port Receive Utilization Percentage measures the approximate utilization percentage of the ports in monitored storage systems. This value represents the average percent of time that the ports were busy receiving data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Port Receive Utilization Percentage value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0608E The Port Send Bandwidth Percentage of port *port name* in device *device name* was measured to be *measured value%*, which violated the critical-stress boundary value of *boundary value%*.

Explanation

Port Send Bandwidth Percentage measures the approximate bandwidth utilization percentage of the ports in monitored storage systems and switches, based on their current negotiated speed. This value represents the ratio of the current send data rate to the maximum send data rate achievable at the negotiated speed. Since fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Port Send Bandwidth Percentage value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this port. Either the port is being overloaded with too high a workload, or there are some congestion or configuration issues in the fabric. If the workload is too high, moving some of the port's workload to other less busy ports might resolve the problem.

If the port's negotiated speed is much less than the physical capability of the port, check the port error metrics which are available for some type of ports, to determine if there are reliability issues with the fabric or the connections. High error counts can indicate hardware problems with some of the ports or switches in the fabric, or can

indicate bad cables or connector between some of the ports.

Also be aware that less capable ports at the subsystem or the server HBA or at the switches along the data path may artificially restrict the speed, so communication can only go as fast as the slowest link on the path. It is therefore usually advisable to have only ports with similar capability connected into a fabric.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0609W The Port Send Bandwidth Percentage of port *port name* in device *device name* was measured to be *measured value%*, which violated the warning-stress boundary value of *boundary value%*.

Explanation

Port Send Bandwidth Percentage measures the approximate bandwidth utilization percentage of the ports in monitored storage systems and switches, based on their current negotiated speed. This value represents the ratio of the current send data rate to the maximum send data rate achievable at the negotiated speed. Since fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Port Send Bandwidth Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this port. Either the port is being overloaded with too high a workload, or there are some congestion or configuration issues in the fabric. If the workload is too high, moving some of the port's workload to other less busy ports might resolve the problem.

If the port's negotiated speed is much less than the physical capability of the port, check the port error metrics which are available for some type of ports, to determine if there are reliability issues with the fabric or the connections. High error counts can indicate hardware problems with some of the ports or switches in the fabric, or can indicate bad cables or connector between some of the ports.

Also be aware that less capable ports at the subsystem or the server HBA or at the switches along the data path may artificially restrict the speed, so communication can only go as fast as the slowest link on the path. It is therefore usually advisable to have only ports with similar capability connected into a fabric.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0610W The Port Send Bandwidth Percentage of port *port name* in device *device name* was measured to be *measured value%*, which violated the warning-idle boundary value of *boundary value%*.

Explanation

Port Send Bandwidth Percentage measures the approximate bandwidth utilization percentage of the ports in monitored storage systems and switches, based on their current negotiated speed. This value represents the ratio of the current send data rate to the maximum send data rate achievable at the negotiated speed. Since fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Port Send Bandwidth Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system or switch are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0611E The Port Send Bandwidth Percentage of port *port name* in device *device name* was measured to be *measured value%*, which violated the critical-idle boundary value of *boundary value%*.

Explanation

Port Send Bandwidth Percentage measures the approximate bandwidth utilization percentage of the ports in monitored storage systems and switches, based on their current negotiated speed. This value represents the ratio of the current send data rate to the maximum send data rate achievable at the negotiated speed. Since fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Port Send Bandwidth Percentage value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system or switch are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0612E The Port Receive Bandwidth Percentage of port *port name* in device *device name* was measured to be *measured value%*, which violated the critical-stress boundary value of *boundary value%*.

Explanation

Port Receive Bandwidth Percentage measures the approximate bandwidth utilization percentage of the ports in monitored storage systems and switches, based on their current negotiated speed. This value represents the ratio of the current receive data rate to the maximum receive data rate achievable at the negotiated speed. Since fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Port Receive Bandwidth Percentage value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this port. Either the port is being overloaded with too high a workload, or there are some congestion or configuration issues in the fabric. If the workload is too high, moving some of the port's workload to other less busy ports might resolve the problem.

If the port's negotiated speed is much less than the physical capability of the port, check the port error metrics which are available for some type of ports, to determine if there are reliability issues with the fabric or the connections. High error counts can indicate hardware problems with some of the ports or switches in the fabric, or can indicate bad cables or connector between some of the ports.

Also be aware that less capable ports at the subsystem or the server HBA or at the switches along the data path may artificially restrict the speed, so communication can only go as fast as the slowest link on the path. It is therefore usually advisable to have only ports with similar capability connected into a fabric.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0613W The Port Receive Bandwidth Percentage of port *port name* in device *device name* was measured to be *measured value%*, which violated the warning-stress boundary value of *boundary value%*.

Explanation

Port Receive Bandwidth Percentage measures the approximate bandwidth utilization percentage of the ports in monitored storage systems and switches, based on their current negotiated speed. This value represents the ratio of the current receive data rate to the maximum receive data rate achievable at the negotiated speed. Since fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Port Receive Bandwidth Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a performance problem exists for this port. Either the port is being overloaded with too high a workload, or there are some congestion or configuration issues in the fabric. If the workload is too high, moving some of the port's workload to other less busy ports might resolve the problem.

If the port's negotiated speed is much less than the physical capability of the port, check the port error metrics which are available for some type of ports, to determine if there are reliability issues with the fabric or the connections. High error counts can indicate hardware problems with some of the ports or switches in the fabric, or can indicate bad cables or connector between some of the ports.

Also be aware that less capable ports at the subsystem or the server HBA or at the switches along the data path may artificially restrict the speed, so communication can only go as fast as the slowest link on the path. It is therefore usually advisable to have only ports with similar capability connected into a fabric.

However this threshold indication is only as good as its defined boundary values. If the specified boundary value is too low, IBM Spectrum Control might recognize many violations, even though there is really no underlying performance problem with the port. In this case, increase the boundary value in the threshold definition to reduce the number of unnecessary alerts.

ALR0614W The Port Receive Bandwidth Percentage of port *port name* in device *device name* was measured to be *measured value*%, which violated the warning-idle boundary value of *boundary value*%.

Explanation

Port Receive Bandwidth Percentage measures the approximate bandwidth utilization percentage of the ports in monitored storage systems and switches, based on their current negotiated speed. This value represents the ratio of the current receive data rate to the maximum receive data rate achievable at the negotiated speed. Since fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Port Receive Bandwidth Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system or switch are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0615E The Port Receive Bandwidth Percentage of port *port name* in device *device name* was measured to be *measured value*%, which violated the critical-idle boundary value of *boundary value*%.

Explanation

Port Receive Bandwidth Percentage measures the approximate bandwidth utilization percentage of the ports in monitored storage systems and switches, based on their current negotiated speed. This value represents the ratio of the current receive data rate to the maximum receive data rate achievable at the negotiated speed. Since fibre channel protocols are full-duplex, the sum of the send and receive bandwidth percentages can theoretically reach a maximum of 200%. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Port Receive Bandwidth Percentage value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

The threshold violation might indicate an unexpected drop in the workload for this port. This can happen if the application(s) using the port stopped working properly.

This type of threshold boundary should only be defined if the workload for the ports in a storage system or switch are guaranteed to always remain at a certain minimum level. Otherwise IBM Spectrum Control might recognize many violations, even though there is really no underlying problem with the applications. In these situations, consider decreasing the boundary value in the threshold definition, or disabling the boundary check by leaving the boundary value blank.

ALR0616E The CRC Error Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

CRC Error Rate measures the average number of frames that were received per second, where the cyclic redundancy check (CRC) value in the frame did not match the CRC value computed by the receiving port. This indicates accidental modification of the frame's data during transit, and can be used as an indicator of the relative quality of the port's connection. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an CRC Error Rate that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the error rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0617W The CRC Error Rate of port *port name* in device *device name* was measured to be *measured value cnt/s*, which violated the defined warning-stress boundary value of *boundary value cnt/s*.

Explanation

CRC Error Rate measures the average number of frames that were received per second, where the cyclic redundancy check (CRC) value in the frame did not match the CRC value computed by the receiving port. This indicates accidental modification of the frame's data during transit, and can be used as an indicator of the relative quality of the port's connection. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an CRC Error Rate that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the error rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0618W The CRC Error Rate of port *port name* in device *device name* was measured to be *measured value cnt/s*, which violated the defined warning-idle boundary value of *boundary value cnt/s*.

Explanation

CRC Error Rate measures the average number of frames that were received per second, where the cyclic redundancy check (CRC) value in the frame did not match the CRC value computed by the receiving port. This indicates accidental modification of the frame's data during transit, and can be used as an indicator of the relative quality of the port's connection. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an CRC Error Rate that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the port error thresholds, because a zero or low error rate is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0619E The CRC Error Rate of port *port name* in device *device name* was measured to be *measured value cnt/s*, which violated the defined critical-idle boundary value of *boundary value cnt/s*.

Explanation

CRC Error Rate measures the average number of frames that were received per second, where the cyclic redundancy check (CRC) value in the frame did not match the CRC value computed by the receiving port. This indicates accidental modification of the frame's data during transit, and can be used as an indicator of the relative quality of the port's connection. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an CRC Error Rate that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the port error thresholds, because a zero or low error rate is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0620E The Invalid Transmission Word Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

Invalid Transmission Word Rate measures the average number of invalid transmission words that were detected by the port, per second. Transmission words are primitive elements used by the fibre channel protocol for transmission of data. A frame consists of between 10 and 537 transmission words. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an Invalid Transmission Word Rate that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the error rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0621W The Invalid Transmission Word Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-stress boundary value of *boundary value* cnt/s.

Explanation

Invalid Transmission Word Rate measures the average number of invalid transmission words that were detected by the port, per second. Transmission words are primitive elements used by the fibre channel protocol for transmission of data. A frame consists of between 10 and 537 transmission words. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an Invalid Transmission Word Rate that is greater than or equal to, the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the error rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0622W The Invalid Transmission Word Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

Invalid Transmission Word Rate measures the average number of invalid transmission words that were detected by the port, per second. Transmission words are primitive elements used by the fibre channel protocol for transmission of data. A frame consists of between 10 and 537 transmission words. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an Invalid Transmission Word Rate that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the port error thresholds, because a zero or low error rate is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0623E The Invalid Transmission Word Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

Invalid Transmission Word Rate measures the average number of invalid transmission words that were detected by the port, per second. Transmission words are primitive elements used by the fibre channel protocol for transmission of data. A frame consists of between 10 and 537 transmission words. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has an Invalid Transmission Word Rate that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the port error thresholds, because a zero or low error rate is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0624E The Zero Buffer Credit Timer of port *port name* in storage system *storage system name* was measured to be *measured value* microseconds, which violated the defined critical-stress boundary value of *boundary value* microseconds.

Explanation

Zero Buffer Credit Timer measures the number of microseconds during which it was impossible to send frames due to a lack of buffer credits at the port. Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Zero Buffer Credit Timer value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the timer value remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0625W The Zero Buffer Credit Timer of port *port name* in storage system *storage system name* was measured to be *measured value* microseconds, which violated the defined warning-stress boundary value of *boundary value* microseconds.

Explanation

Zero Buffer Credit Timer measures the number of microseconds during which it was impossible to send frames due to a lack of buffer credits at the port. Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Zero Buffer Credit Timer value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the timer value remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0626W The Zero Buffer Credit Timer of port *port name* in storage system *storage system name* was measured to be *measured value* microseconds, which violated the defined warning-idle boundary value of *boundary value* microseconds.

Explanation

Zero Buffer Credit Timer measures the number of microseconds during which it was impossible to send frames due to a lack of buffer credits at the port. Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Zero Buffer Credit Timer value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Zero Buffer Credit Timer threshold, because a zero or low time value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0627E The Zero Buffer Credit Timer of port *port name* in storage system *storage system name* was measured to be *measured value* microseconds, which violated the defined critical-idle boundary value of *boundary value* microseconds.

Explanation

Zero Buffer Credit Timer measures the number of microseconds during which it was impossible to send frames due to a lack of buffer credits at the port. Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system or switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Zero Buffer Credit Timer value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Zero Buffer Credit Timer threshold, because a zero or low time value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0628E The Zero Buffer Credit Percentage of port *port name* in storage system *storage system name* was measured to be *measured value*%, which violated the defined critical-stress boundary value of *boundary value*%.

Explanation

Zero Buffer Credit Percentage measures the percentage of time that the port was unable to send frames due to a lack of buffer credits at the port over a particular time interval. Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Zero Buffer Credit Percentage value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the timer value remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the storage system, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0629W The Zero Buffer Credit Percentage of port *port name* in storage system *storage system name* was measured to be *measured value%*, which violated the defined warning-stress boundary value of *boundary value%*.

Explanation

Zero Buffer Credit Percentage measures the percentage of time that the port was unable to send frames due to a lack of buffer credits at the port over a particular time interval. Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Zero Buffer Credit Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the timer value remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the storage system, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0630W The Zero Buffer Credit Percentage of port *port name* in storage system *storage system name* was measured to be *measured value%*, which violated the defined warning-idle boundary value of *boundary value%*.

Explanation

Zero Buffer Credit Percentage measures the percentage of time that the port was unable to send frames due to a lack of buffer credits at the port over a particular time interval. Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Zero Buffer Credit Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Zero Buffer Credit Percentage threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0631E The Zero Buffer Credit Percentage of port *port name* in storage system *storage system name* was measured to be *measured value%*, which violated the defined critical-idle boundary value of *boundary value%*.

Explanation

Zero Buffer Credit Percentage measures the percentage of time that the port was unable to send frames due to a lack of buffer credits at the port over a particular time interval. Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not

overrun the ability of the receiving port to receive that data. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system has a Zero Buffer Credit Percentage value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Zero Buffer Credit Percentage threshold, because a zero or low time value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0632E The Discarded Frame Rate of port *port name* in switch *switch name* was measured to be *measured value%*, which violated the defined critical-stress boundary value of *boundary value%*.

Explanation

The Discarded Frame Rate measures the rate at which the port discarded frames for any reason. Frames can be discarded due to many reasons, for example invalid addressing or header contents, or due to timeout caused by lack of buffer credits. Other error metrics may help in narrowing down the reason for the discards. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Discarded Frame Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0633W The Discarded Frame Rate of port *port name* in switch *switch name* was measured to be *measured value%*, which violated the defined warning-stress boundary value of *boundary value%*.

Explanation

The Discarded Frame Rate measures the rate at which the port discarded frames for any reason. Frames can be discarded due to many reasons, for example invalid addressing or header contents, or due to timeout caused by lack of buffer credits. Other error metrics may help in narrowing down the reason for the discards. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Discarded Frame Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0634W The Discarded Frame Rate of port *port name* in switch *switch name* was measured to be *measured value%*, which violated the defined warning-idle boundary value of *boundary value%*.

Explanation

The Discarded Frame Rate measures the rate at which the port discarded frames for any reason. Frames can be discarded due to many reasons, for example invalid addressing or header contents, or due to timeout caused by lack of buffer credits. Other error metrics may help in narrowing down the reason for the discards. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Discarded Frame Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Discarded Frame Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0635E The Discarded Frame Rate of port *port name* in switch *switch name* was measured to be *measured value%*, which violated the defined critical-idle boundary value of *boundary value%*.

Explanation

The Discarded Frame Rate measures the rate at which the port discarded frames for any reason. Frames can be discarded due to many reasons, for example invalid addressing or header contents, or due to timeout caused by lack of buffer credits. Other error metrics may help in narrowing down the reason for the discards. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Discarded Frame Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Discarded Frame Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0636E The Loss of Sync Rate of port *port name* in device *device name* was measured to be *measured value cnt/s*, which violated the defined critical-stress boundary value of *boundary value cnt/s*.

Explanation

The Loss of Sync Rate measures the rate at which the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Loss of Sync Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected port. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0637W The Loss of Sync Rate of port *port name* in device *device name* was measured to be *measured value cnt/s*, which violated the defined warning-stress boundary value of *boundary value cnt/s*.

Explanation

The Loss of Sync Rate measures the rate at which the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Loss of Sync Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected port. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0638W The Loss of Sync Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

The Loss of Sync Rate measures the rate at which the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Loss of Sync Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Loss of Sync Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0639E The Loss of Sync Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

The Loss of Sync Rate measures the rate at which the port lost synchronization with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to mismatching port speeds between the partner ports, when auto-negotiation of link speed is disabled. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Loss of Sync Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Loss of Sync Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0640E The Loss of Signal Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

The Loss of Signal Rate measures the rate at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to exceeding the maximum link distance between ports, for the type of connecting cable and light source. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Loss of Signal Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware problem exists either for this port or for the connected port. However it could also simply mean that the connected port was reset somehow, for example by rebooting the host containing the HBA. Therefore, isolated errors should be ignored in most cases, but if the rate remains

consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0641W The Loss of Signal Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-stress boundary value of *boundary value* cnt/s.

Explanation

The Loss of Signal Rate measures the rate at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to exceeding the maximum link distance between ports, for the type of connecting cable and light source. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Loss of Signal Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware problem exists either for this port or for the connected port. However it could also simply mean that the connected port was reset somehow, for example by rebooting the host containing the HBA. Therefore, isolated errors should be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0642W The Loss of Signal Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

The Loss of Signal Rate measures the rate at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to exceeding the maximum link distance between ports, for the type of connecting cable and light source. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Loss of Signal Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Loss of Signal Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0643E The Loss of Signal Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

The Loss of Signal Rate measures the rate at which the port lost communication with its partner port. These types of errors usually indicate physical link problems, caused by faulty SFP modules or cables, or caused by faulty connections at the switch or patch panel. However in some cases this can also occur due to exceeding the maximum link distance between ports, for the type of connecting cable and light source. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Loss of Signal Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Loss of Signal Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0644E The Discarded Class 3 Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

The Discarded Class 3 Frame Rate measures the rate at which the port Discarded Class 3 Frames for any reason. Frames can be discarded due to many reasons, for example invalid addressing or header contents, or due to timeout caused by lack of buffer credits. Other error metrics may help in narrowing down the reason for the discards. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Discarded Class 3 Frame Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0645W The Discarded Class 3 Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined warning-stress boundary value of *boundary value* cnt/s.

Explanation

The Discarded Class 3 Frame Rate measures the rate at which the port Discarded Class 3 Frames for any reason. Frames can be discarded due to many reasons, for example invalid addressing or header contents, or due to timeout caused by lack of buffer credits. Other error metrics may help in narrowing down the reason for the discards. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Discarded Class 3 Frame Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0646W The Discarded Class 3 Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

The Discarded Class 3 Frame Rate measures the rate at which the port Discarded Class 3 Frames for any reason. Frames can be discarded due to many reasons, for example invalid addressing or header contents, or due to timeout caused by lack of buffer credits. Other error metrics may help in narrowing down the reason for the discards. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Discarded Class 3 Frame Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Discarded Class 3 Frame Rate threshold, because a zero or low time value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0647E The Discarded Class 3 Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

The Discarded Class 3 Frame Rate measures the rate at which the port Discarded Class 3 Frames for any reason. Frames can be discarded due to many reasons, for example invalid addressing or header contents, or due to timeout caused by lack of buffer credits. Other error metrics may help in narrowing down the reason for the discards. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Discarded Class 3 Frame Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Discarded Class 3 Frame Rate threshold, because a zero or low time value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0648E The Zero Buffer Credit Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

Zero Buffer Credit Rate measures the percentage of time that the port was unable to send frames due to a lack of buffer credits at the port over a particular time interval. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the port's overall throughput, and can cascade from port to port causing congestion in the entire fabric. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Zero Buffer Credit Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0649W The Zero Buffer Credit Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined warning-stress boundary value of *boundary value* cnt/s.

Explanation

Zero Buffer Credit Rate measures the percentage of time that the port was unable to send frames due to a lack of buffer credits at the port over a particular time interval. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the port's overall throughput, and can cascade from port to port causing congestion in the entire fabric. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Zero Buffer Credit Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0650W The Zero Buffer Credit Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

Zero Buffer Credit Rate measures the percentage of time that the port was unable to send frames due to a lack of buffer credits at the port over a particular time interval. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the port's overall throughput, and can cascade from port to port causing congestion in the entire fabric. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Zero Buffer Credit Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Zero Buffer Credit Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0651E The Zero Buffer Credit Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

Zero Buffer Credit Rate measures the percentage of time that the port was unable to send frames due to a lack of buffer credits at the port over a particular time interval. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the port's overall throughput, and can cascade from port to port causing congestion in the entire fabric. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Zero Buffer Credit Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Zero Buffer Credit Rate threshold, because a zero or low time value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0652E The Class 3 Send Timeout Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s. This

violated the defined critical-stress boundary value of *boundary value cnt/s*.

Explanation

Class 3 Send Timeout Frame Rate measures the rate at which class 3 frames had to be discarded at the transmission port due to timeout caused by lack of buffer credits. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. When out of buffer credits, the port waits for a configurable amount of time (e.g. 500 ms) to receive credit and send the waiting frame. If no credit is received, a timeout condition occurs and the waiting frame is discarded. Note that if egress and ingress ports are on the same blade, the Class 3 Receive Timeout Frame rate increases for the ingress port(s) at the same time the Class 3 Send Timeout Frame rate increases for the egress port.

A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch. This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Class 3 Send Timeout Frame Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0653W The Class 3 Send Timeout Frame Rate of port *port name* in switch *switch name* was measured to be *measured value cnt/s*. This violated the defined warning-stress boundary value of *boundary value cnt/s*.

Explanation

Class 3 Send Timeout Frame Rate measures the rate at which class 3 frames had to be discarded at the transmission port due to timeout caused by lack of buffer credits. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. When out of buffer credits, the port waits for a configurable amount of time (e.g. 500 ms) to receive credit and send the waiting frame. If no credit is received, a timeout condition occurs and the waiting frame is discarded. Note that if egress and ingress ports are on the same blade, the Class 3 Receive Timeout Frame rate increases for the ingress port(s) at the same time the Class 3 Send Timeout Frame rate increases for the egress port.

A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch. This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Class 3 Send Timeout Frame Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0654W The Class 3 Send Timeout Frame Rate of port *port name* in switch *switch name* was measured to be *measured value cnt/s*. This violated the defined warning-idle boundary value of *boundary value cnt/s*.

Explanation

Class 3 Send Timeout Frame Rate measures the rate at which class 3 frames had to be discarded at the transmission port due to timeout caused by lack of buffer credits. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. When out of buffer credits, the port waits for a configurable amount of time (e.g. 500 ms) to receive credit and send the waiting frame. If no credit is received, a timeout condition occurs and the waiting frame is discarded. Note that if egress and ingress ports are on the same blade, the Class 3 Receive Timeout Frame rate increases for the ingress port(s) at the same time the Class 3 Send Timeout Frame rate increases for the egress port.

A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch. This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Class 3 Send Timeout Frame Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Class 3 Send Timeout Frame Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0655E The Class 3 Send Timeout Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s. This violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

Class 3 Send Timeout Frame Rate measures the rate at which class 3 frames had to be discarded at the transmission port due to timeout caused by lack of buffer credits. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. When out of buffer credits, the port waits for a configurable amount of time (e.g. 500 ms) to receive credit and send the waiting frame. If no credit is received, a timeout condition occurs and the waiting frame is discarded. Note that if egress and ingress ports are on the same blade, the Class 3 Receive Timeout Frame rate increases for the ingress port(s) at the same time the Class 3 Send Timeout Frame rate increases for the egress port.

A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch. This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Class 3 Send Timeout Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Class 3 Send Timeout Frame Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0656E The Class 3 Receive Timeout Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s. This violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

Class 3 Receive Timeout Frame Rate measures the rate at which class 3 frames had to be discarded at the receiving port due to timeout caused by lack of buffer credits at the transmitting port. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. When out of buffer credits, the port waits for a configurable amount of time (e.g. 500 ms) to receive credit and send the waiting frame. If no credit is received, a timeout condition occurs and the waiting frame is discarded. Note that if egress and ingress ports are on the same blade, the Class 3 Receive Timeout Frame rate increases for the ingress port(s) at the same time the Class 3 Send Timeout Frame rate increases for the egress port.

A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch. This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Class 3 Receive Timeout Frame Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0657W The Class 3 Receive Timeout Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s. This violated the defined warning-stress boundary value of *boundary value* cnt/s.

Explanation

Class 3 Receive Timeout Frame Rate measures the rate at which class 3 frames had to be discarded at the receiving port due to timeout caused by lack of buffer credits at the transmitting port. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is

sending data does not overrun the ability of the receiving port to receive that data. When out of buffer credits, the port waits for a configurable amount of time (e.g. 500 ms) to receive credit and send the waiting frame. If no credit is received, a timeout condition occurs and the waiting frame is discarded. Note that if egress and ingress ports are on the same blade, the Class 3 Receive Timeout Frame rate increases for the ingress port(s) at the same time the Class 3 Send Timeout Frame rate increases for the egress port.

A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch. This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Class 3 Receive Timeout Frame Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0658W The Class 3 Receive Timeout Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s. This violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

Class 3 Receive Timeout Frame Rate measures the rate at which class 3 frames had to be discarded at the receiving port due to timeout caused by lack of buffer credits at the transmitting port. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. When out of buffer credits, the port waits for a configurable amount of time (e.g. 500 ms) to receive credit and send the waiting frame. If no credit is received, a timeout condition occurs and the waiting frame is discarded. Note that if egress and ingress ports are on the same blade, the Class 3 Receive Timeout Frame rate increases for the ingress port(s) at the same time the Class 3 Send Timeout Frame rate increases for the egress port.

A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch. This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Class 3 Receive Timeout Frame Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Class 3 Receive Timeout Frame Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0659E The Class 3 Receive Timeout Frame Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s. This violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

Class 3 Receive Timeout Frame Rate measures the rate at which class 3 frames had to be discarded at the receiving port due to timeout caused by lack of buffer credits at the transmitting port. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port which is sending data does not overrun the ability of the receiving port to receive that data. When out of buffer credits, the port waits for a configurable amount of time (e.g. 500 ms) to receive credit and send the waiting frame. If no credit is received, a timeout condition occurs and the waiting frame is discarded. Note that if egress and ingress ports are on the same blade, the Class 3 Receive Timeout Frame rate increases for the ingress port(s) at the same time the Class 3 Send Timeout Frame rate increases for the egress port.

A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch. This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Class 3 Receive Timeout Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Class 3 Receive Timeout Frame Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0660E The Credit Recovery Link Reset Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which

violated the defined critical-stress boundary value of *boundary value cnt/s*.

Explanation

Credit Recovery Link Reset Rate estimates the rate at which link resets were performed in order to recover buffer credits. Link resets may be performed for other reasons, but this metric estimates the rate for only buffer credit recovery by ignoring any resets that may have been caused by loss of synchronization or loss of signal conditions. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Credit Recovery Link Reset Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0661W The Credit Recovery Link Reset Rate of port *port name* in device *device name* was measured to be *measured value cnt/s*, which violated the defined warning-stress boundary value of *boundary value cnt/s*.

Explanation

Credit Recovery Link Reset Rate estimates the rate at which link resets were performed in order to recover buffer credits. Link resets may be performed for other reasons, but this metric estimates the rate for only buffer credit recovery by ignoring any resets that may have been caused by loss of synchronization or loss of signal conditions. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Credit Recovery Link Reset Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0662W The Credit Recovery Link Reset Rate of port *port name* in device *device name* was measured to be *measured value cnt/s*, which violated the defined warning-idle boundary value of *boundary value cnt/s*.

Explanation

Credit Recovery Link Reset Rate estimates the rate at which link resets were performed in order to recover buffer credits. Link resets may be performed for other reasons, but this metric estimates the rate for only buffer credit recovery by ignoring any resets that may have been caused by loss of synchronization or loss of signal conditions. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Credit Recovery Link Reset Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Credit Recovery Link Reset Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0663E The Credit Recovery Link Reset Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

Credit Recovery Link Reset Rate estimates the rate at which link resets were performed in order to recover buffer credits. Link resets may be performed for other reasons, but this metric estimates the rate for only buffer credit recovery by ignoring any resets that may have been caused by loss of synchronization or loss of signal conditions. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Credit Recovery Link Reset Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Credit Recovery Link Reset Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0664E The RDY Priority Override Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

RDY Priority Override Rate measures the rate at which sending R_RDY or VC_RDY signals was a higher priority for the port than sending frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a RDY Priority Override Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0665W The RDY Priority Override Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined warning-stress boundary value of *boundary value* cnt/s.

Explanation

RDY Priority Override Rate measures the rate at which sending R_RDY or VC_RDY signals was a higher priority for the port than sending frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a RDY Priority Override Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0666W The RDY Priority Override Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

RDY Priority Override Rate measures the rate at which sending R_RDY or VC_RDY signals was a higher priority for the port than sending frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a RDY Priority Override Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the RDY Priority Override Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0667E The RDY Priority Override Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

RDY Priority Override Rate measures the rate at which sending R_RDY or VC_RDY signals was a higher priority for the port than sending frames. This condition occurs because of diminishing credit reserves in the transmitter at the other end of the fibre. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a RDY Priority Override Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the RDY Priority Override Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0668E The Port State Change Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

Port State Change Rate measures the rate at which the port has changed state between online, offline, and faulty. This threshold can be used to help to identify flapping ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Port State Change Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0669W The Port State Change Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined warning-stress boundary value of *boundary value* cnt/s.

Explanation

Port State Change Rate measures the rate at which the port has changed state between online, offline, and faulty. This threshold can be used to help to identify flapping ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Port State Change Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0670W The Port State Change Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

Port State Change Rate measures the rate at which the port has changed state between online, offline, and faulty. This threshold can be used to help to identify flapping ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Port State Change Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Port State Change Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0671E The Port State Change Rate of port *port name* in switch *switch name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

Port State Change Rate measures the rate at which the port has changed state between online, offline, and faulty. This threshold can be used to help to identify flapping ports. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the switch has a Port State Change Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Port State Change Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0672E The Primitive Sequence Protocol Error Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

Primitive Sequence Protocol Error Rate measures the rate at which the port received unexpected primitive sequences. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Primitive Sequence Protocol Error Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0673W The Primitive Sequence Protocol Error Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-stress boundary value of *boundary value* cnt/s.

Explanation

Primitive Sequence Protocol Error Rate measures the rate at which the port received unexpected primitive sequences. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Primitive Sequence Protocol Error Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0674W The Primitive Sequence Protocol Error Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

Primitive Sequence Protocol Error Rate measures the rate at which the port received unexpected primitive sequences. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Primitive Sequence Protocol Error Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Primitive Sequence Protocol Error Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0675E The Primitive Sequence Protocol Error Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

Primitive Sequence Protocol Error Rate measures the rate at which the port received unexpected primitive sequences. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Primitive Sequence Protocol Error Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Primitive Sequence Protocol Error Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0676E The Link Reset Transmitted Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

Link Reset Transmitted Rate measures the rate at which the port encountered link resets. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Reset Transmitted Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0677W The Link Reset Transmitted Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-stress boundary value of *boundary value* cnt/s.

Explanation

Link Reset Transmitted Rate measures the rate at which the port encountered link resets. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Reset Transmitted Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0678W The Link Reset Transmitted Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined warning-idle boundary value of *boundary value* cnt/s.

Explanation

Link Reset Transmitted Rate measures the rate at which the port encountered link resets. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Reset Transmitted Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Link Reset Transmitted Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0679E The Link Reset Transmitted Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-idle boundary value of *boundary value* cnt/s.

Explanation

Link Reset Transmitted Rate measures the rate at which the port encountered link resets. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Reset Transmitted Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Link Reset Transmitted Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0680E The Link Reset Received Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which violated the defined critical-stress boundary value of *boundary value* cnt/s.

Explanation

Link Reset Received Rate measures the rate at which the port encountered link resets. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Reset Received Rate value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0681W The Link Reset Received Rate of port *port name* in device *device name* was measured to be *measured value* cnt/s, which

violated the defined warning-stress boundary value of *boundary value cnt/s*.

Explanation

Link Reset Received Rate measures the rate at which the port encountered link resets. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Reset Received Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

Action

The threshold violation might indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can be ignored in most cases, but if the rate remains consistently high over time, this may be a cause for concern. Follow the trouble-shooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can be ignored in most cases, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can greatly reduce the number of unnecessary alerts.

ALR0682W The Link Reset Received Rate of port *port name* in device *device name* was measured to be *measured value cnt/s*, which violated the defined warning-idle boundary value of *boundary value cnt/s*.

Explanation

Link Reset Received Rate measures the rate at which the port encountered link resets. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Reset Received Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Link Reset Received Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0683E The Link Reset Received Rate of port *port name* in device *device name* was measured to be *measured value cnt/s*, which violated the defined critical-idle boundary value of *boundary value cnt/s*.

Explanation

Link Reset Received Rate measures the rate at which the port encountered link resets. A threshold was defined on this metric, which causes the measured value to be compared to the defined boundaries for each set of performance statistics collected from the switch.

This message indicates that a threshold boundary violation occurred. In this case, the specified port in the storage system or switch has a Link Reset Received Rate value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

It makes little sense to enable an idle boundary for the Link Reset Received Rate threshold, because a zero or low value is the ideal and expected situation. It is recommended to disable the idle boundary checking by leaving the idle boundary values blank in the related threshold definition.

ALR0684E The Port Congestion Index of port *port_name* in device *device_name* was found to be *measured_value* Counts, which violated the defined critical-stress boundary value of *boundary_value* Counts.

Explanation

Port Congestion Index estimates the degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally between 0 (no congestion) to 100, but can exceed 100 if buffer credit exhaustion persisted for an extended amount of time. A threshold is defined on this metric, which can be used to compare the value of Port Congestion Index with the defined boundaries for each set of performance statistics collected from the storage system or switch.

A threshold boundary violation occurred. The specified port in the storage system or switch has a Port Congestion Index value that is greater than or equal to the critical-stress boundary that was defined for the related threshold.

The threshold violation may indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can generally be ignored, but if the rate remains consistently high, the errors may be a cause for concern.

Action

Contact your resource vendor for troubleshooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

ALR0685W The Port Congestion Index of port *port_name* in device *device_name* was found to be *measured_value* Counts, which violated the defined warning-stress boundary value of *boundary_value* Counts.

Explanation

Port Congestion Index estimates the degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally between 0 (no congestion) to 100, but can exceed 100 if buffer credit exhaustion persisted for an extended amount of time. A threshold is defined on this metric, which can be used to compare the value of Port Congestion Index with the defined boundaries for each set of performance statistics collected from the storage system or switch.

A threshold boundary violation occurred. The specified port in the storage system or switch has a Port Congestion Index value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary that was defined for the related threshold.

The threshold violation may indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can generally be ignored, but if the rate remains consistently high, the errors may be a cause for concern.

Action

Contact your resource vendor for troubleshooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

ALR0686W The Port Congestion Index of port *port_name* in device *device_name* was found to be *measured_value* Counts, which violated the defined warning-idle boundary value of *boundary_value* Counts.

Explanation

Port Congestion Index estimates the degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally between 0 (no congestion) to 100, but can exceed 100 if buffer credit exhaustion persisted for an extended amount of time. A threshold is defined on this metric, which can be used to compare the value of Port Congestion Index with the defined boundaries for each set of performance statistics collected from the storage system or switch.

A threshold boundary violation occurred. The specified port in the storage system or switch has a Port Congestion Index value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary that was defined for the related threshold.

Action

A zero or low Port Congestion Index value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0687E The Port Congestion Index of port *port_name* in device *device_name* was found to be *measured_value* Counts, which violated the defined critical-idle boundary value of *boundary_value* Counts.

Explanation

Port Congestion Index estimates the degree to which frame transmission was delayed due to a lack of buffer credits. This value is generally between 0 (no congestion) to 100, but can exceed 100 if buffer credit exhaustion persisted for an extended amount of time. A threshold is defined on this metric, which can be used to compare the value of Port Congestion Index with the defined boundaries for each set of performance statistics collected from the storage system or switch.

A threshold boundary violation occurred. The specified port in the storage system or switch has a Port Congestion Index value that is less than or equal to the critical-idle boundary that was defined for the related threshold.

Action

A zero or low Port Congestion Index value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0688E The Link Quality Percentage of port *port_name* in device *device_name* was found to be *measured_value*, which violated the defined critical-stress boundary value of *boundary_value*.

Explanation

Link Quality Percentage estimates the relative quality of the connection of a switch port. The percentage is based on whether the port is an expansion port (E_port) or a fabric port (F_port), and on the numbers and types of errors that are detected by the port. A threshold is defined on this metric, which can be used to compare the value of the Link Quality Percentage with the defined boundaries for each set of performance statistics collected from the switch.

A threshold boundary violation occurred. The specified switch port has a Link Quality Percentage value that is less than or equal to the critical-stress boundary defined for the related threshold.

The threshold violation may indicate that a hardware or configuration problem exists either for this port or for the connected port or cable. Isolated errors can generally be ignored, but if the percentage remains consistently high, the errors may be a cause for concern.

Action

Contact your resource vendor for troubleshooting guidelines for the switch, and for the connected resource (storage system, switch, or server HBA).

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

ALR0689W The Link Quality Percentage of port *port_name* in device *device_name* was found to be *measured_value*, which violated the defined warning-stress boundary value of *boundary_value*.

Explanation

Link Quality Percentage estimates the relative quality of the connection of a switch port. The percentage is based on whether the port is an expansion port (E_port) or a fabric port (F_port), and on the numbers and types of errors that are detected by the port. A threshold is defined on this metric, which can be used to compare the value of the Link Quality Percentage with the defined boundaries for each set of performance statistics collected from the switch.

A threshold boundary violation occurred. The specified switch port has a Link Quality Percentage value that is less than or equal to the warning-stress boundary, but greater than the critical-stress boundary defined for the related threshold.

The threshold violation may indicate that a hardware or configuration problem exists either for this port or for the connected port or cable. Isolated errors can generally be ignored, but if the percentage remains consistently high, the errors may be a cause for concern.

Action

Contact your resource vendor for troubleshooting guidelines for the switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

ALR0690W The Link Quality Percentage of port *port_name* in device *device_name* was found to be *measured_value*, which violated the defined warning-idle boundary value of *boundary_value*.

Explanation

Link Quality Percentage estimates the relative quality of the connection of a switch port. The percentage is based on whether the port is an expansion port (E_port) or a fabric port (F_port), and on the numbers and types of errors that are detected by the port. A threshold is defined on this metric, which can be used to compare the value of the Link Quality Percentage with the defined boundaries for each set of performance statistics collected from the switch.

A threshold boundary violation occurred. The specified switch port has a Link Quality Percentage value that is greater than or equal to the warning-idle boundary, but less than the critical-idle boundary defined for the related threshold.

Action

A high Link Quality Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0691E The Link Quality Percentage of port *port_name* in device *device_name* was found to be *measured_value*, which violated the defined critical-idle boundary value of *boundary_value*.

Explanation

Link Quality Percentage estimates the relative quality of the connection of a switch port. The percentage is based on whether the port is an expansion port (E_port) or a fabric port (F_port), and on the numbers and types of errors that are detected by the port. A threshold is defined on this metric, which can be used to compare the value of the Link Quality Percentage with the defined boundaries for each set of performance statistics collected from the switch.

A threshold boundary violation occurred. The specified switch port has a Link Quality Percentage value that is greater than or equal to the critical-idle boundary defined for the related threshold.

Action

A high Link Quality Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0692E The Invalid Link Transmission Rate of port *port_name* in device *device_name* was found to be *measured_value*, which violated the defined critical-stress boundary value of *boundary_value*.

Explanation

Invalid Link Transmission Rate measures the average number of invalid transmission words that were detected by the port, per second, in the absence of any synchronization or signal loss on the link. This is the rate of invalid transmission words during time intervals when the link is otherwise up and operating normally. Transmission words are primitive elements used by the fibre channel protocol for transmission of data. A frame consists of between 10 and 537 transmission words. A threshold is defined on this metric, which can be used to compare the value of Invalid Link Transmission Rate with the defined boundaries for each set of performance statistics collected from the storage system or switch.

A threshold boundary violation occurred. The specified storage system or switch port has an Invalid Link Transmission Rate value that is greater than or equal to the critical-stress boundary defined for the related threshold.

The threshold violation may indicate that a hardware or configuration problem exists either for this port or for the connected port or cable. Isolated errors can generally be ignored, but if the percentage remains consistently high, the errors may be a cause for concern.

Action

Contact your resource vendor for troubleshooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

ALR0693W The Invalid Link Transmission Rate of port *port_name* in device *device_name* was found to be *measured_value*, which violated the defined warning-stress boundary value of *boundary_value*.

Explanation

Invalid Link Transmission Rate measures the average number of invalid transmission words that were detected by the port, per second, in the absence of any synchronization or signal loss on the link. This is the rate of invalid transmission words during time intervals when the link is otherwise up and operating normally. Transmission words are primitive elements used by the fibre channel protocol for transmission of data. A frame consists of between 10 and 537 transmission words. A

threshold is defined on this metric, which can be used to compare the value of Invalid Link Transmission Rate with the defined boundaries for each set of performance statistics collected from the storage system or switch.

A threshold boundary violation occurred. The specified storage system or switch port has an Invalid Link Transmission Rate value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary defined for the related threshold.

The threshold violation may indicate that a hardware or configuration problem exists either for this port or for the connected port or cable. Isolated errors can generally be ignored, but if the percentage remains consistently high, the errors may be a cause for concern.

Action

Contact your resource vendor for troubleshooting guidelines for the storage system or switch, and for the connected entity (storage system, switch, or server HBA).

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

ALR0694W The Invalid Link Transmission Rate of port *port_name* in device *device_name* was found to be *measured_value*, which violated the defined warning-idle boundary value of *boundary_value*.

Explanation

Invalid Link Transmission Rate measures the average number of invalid transmission words that were detected by the port, per second, in the absence of any synchronization or signal loss on the link. This is the rate of invalid transmission words during time intervals when the link is otherwise up and operating normally. Transmission words are primitive elements used by the fibre channel protocol for transmission of data. A frame consists of between 10 and 537 transmission words. A threshold is defined on this metric, which can be used to compare the value of Invalid Link Transmission Rate with the defined boundaries for each set of performance statistics collected from the storage system or switch.

A threshold boundary violation occurred. The specified storage system or switch port has an Invalid Link Transmission Rate value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary defined for the related threshold.

Action

A zero or low Invalid Link Transmission Rate value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0695E The Invalid Link Transmission Rate of port *port_name* in device *device_name* was found to be *measured_value*, which violated the defined critical-idle boundary value of *boundary_value*.

Explanation

Invalid Link Transmission Rate measures the average number of invalid transmission words that were detected by the port, per second, in the absence of any synchronization or signal loss on the link. This is the rate of invalid transmission words during time intervals when the link is otherwise up and operating normally. Transmission words are primitive elements used by the fibre channel protocol for transmission of data. A frame consists of between 10 and 537 transmission words. A threshold is defined on this metric, which can be used to compare the value of Invalid Link Transmission Rate with the defined boundaries for each set of performance statistics collected from the storage system or switch.

A threshold boundary violation occurred. The specified storage system or switch port has an Invalid Link Transmission Rate value that is less than or equal to the critical-idle boundary defined for the related threshold.

Action

A zero or low Invalid Link Transmission Rate value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0696E The Extreme I/O Concurrency Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined critical-stress boundary value of *boundary_value*.

Explanation

Extreme I/O Concurrency Percentage measures the percentage of I/O operations that were issued against the port when it already had a high number of outstanding open exchanges pending. This metric shows how often there was too much concurrent I/O occurring on the port. The number of concurrent exchanges for a port on a DS8000

storage system cannot exceed 2000. Approximately 1500 open exchanges for port is a high number for a DS8000 system. A threshold is defined on this metric, which can be used to compare the value of the Extreme I/O Concurrency Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an Extreme I/O Concurrency Percentage value that is greater than or equal to the critical-stress boundary defined for the related threshold.

Action

The threshold violation indicates that there is a danger that the physical capability of the port on the DS8000 system might be exceeded. If the port reaches its maximum number of open exchanges, it responds to any additional I/O requests with a 'Busy' status. This usually means that the requesting host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

If the Extreme I/O Concurrency Percentage remains high for extended periods of time, check the workload of the affected DS8000 system to ensure that it is generally balanced across all the HA ports. If the workload is not balanced across all the ports, try to balance it. Ensure that multipath device drivers are installed and functioning properly on the host servers, so that I/O activity can be spread across multiple DS8000 ports.

If multipath drivers are functioning normally, consider configuring certain ports to allow access from only certain host servers, or making zoning changes between host and storage system ports. Sometimes, you might need to change the cabling or the fabric configuration.

If the workload is already balanced and the Extreme I/O Concurrency Percentage remains high, consider installing more ports in the DS8000 system, or consider moving some of the workload to other storage systems.

ALR0697W The Extreme I/O Concurrency Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined warning-stress boundary value of *boundary_value*.

Explanation

Extreme I/O Concurrency Percentage measures the percentage of I/O operations that were issued against the port when it already had a high number of outstanding open exchanges pending. This metric shows how often there was too much concurrent I/O occurring on the port. The number of concurrent exchanges for a port on a DS8000 storage system cannot exceed 2000. Approximately 1500 open exchanges for port is a high number for a DS8000 system. A threshold is defined on this metric, which can be used to compare the value of the Extreme I/O Concurrency Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an Extreme I/O Concurrency Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary defined for the related threshold.

Action

The threshold violation indicates that there is a danger that the physical capability of the port on the DS8000 system might be exceeded. If the port reaches its maximum number of open exchanges, it responds to any additional I/O requests with a BUSY status. This usually means that the requesting host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

If the Extreme I/O Concurrency Percentage remains high for extended periods of time, check the workload of the affected DS8000 system to ensure that it is generally balanced across all the HA ports. If the workload is not balanced across all the ports, try to balance it. Ensure that multipath device drivers are installed and functioning properly on the host servers, so that I/O activity can be spread across multiple DS8000 ports.

If multipath drivers are functioning normally, consider configuring certain ports to allow access from only certain host servers, or making zoning changes between host and storage system ports. Sometimes, you might need to change the cabling or the fabric configuration.

If the workload is already balanced and the Extreme I/O Concurrency Percentage remains high, consider installing more ports in the DS8000 system, or consider moving some of the workload to other storage systems.

ALR0698W The Extreme I/O Concurrency Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined warning-idle boundary value of *boundary_value*.

Explanation

Extreme I/O Concurrency Percentage measures the percentage of I/O operations that were issued against the port when it already had a high number of outstanding open exchanges pending. This metric shows how often there was too much concurrent I/O occurring on the port. The number of concurrent exchanges for a port on a DS8000

storage system cannot exceed 2000. Approximately 1500 open exchanges for port is a high number for a DS8000 system. A threshold is defined on this metric, which can be used to compare the value of the Extreme I/O Concurrency Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an Extreme I/O Concurrency Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary defined for the related threshold.

Action

A zero or low Extreme I/O Concurrency Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0699E The Extreme I/O Concurrency Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined critical-idle boundary value of *boundary_value*.

Explanation

Extreme I/O Concurrency Percentage measures the percentage of I/O operations that were issued against the port when it already had a high number of outstanding open exchanges pending. This metric shows how often there was too much concurrent I/O occurring on the port. The number of concurrent exchanges for a port on a DS8000 storage system cannot exceed 2000. Approximately 1500 open exchanges for port is a high number for a DS8000 system. A threshold is defined on this metric, which can be used to compare the value of the Extreme I/O Concurrency Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an Extreme I/O Concurrency Percentage value that is less than or equal to the critical-idle boundary defined for the related threshold.

Action

A zero or low Extreme I/O Concurrency Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0700E The I/O Busy Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined critical-stress boundary value of *boundary_value*.

Explanation

I/O Busy Percentage measures the percentage of I/O operations for which the port returned a 'SCSI Queue Full' or 'Busy' status to the host server. These statuses can occur for ports on DS8000 storage systems when the number of open exchanges for the port has reached the maximum value of 2000. This situation can happen when there is too much concurrent I/O being driven to the port. A threshold is defined on this metric, which causes the value of the I/O Busy Percentage to be compared to the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an I/O Busy Percentage value that is greater than or equal to the critical-stress boundary defined for the related threshold.

Action

The threshold violation indicates that the physical capability of the port on the DS8000 system is being exceeded. When a 'Busy' status is returned by the port, this usually means that the requesting host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

If the I/O Busy Percentage remains high for extended periods of time, check the workload of the affected DS8000 system to ensure that it is generally balanced across all the HA ports. If the workload is not balanced across all the ports, try to balance it. Ensure that multipath device drivers are installed and functioning properly on the host servers, so that I/O activity can be spread across multiple DS8000 ports.

If multipath drivers are functioning normally, consider configuring certain ports to allow access from only certain host servers, or consider making zoning changes between host and storage system ports. Sometimes, you might need to change the cabling or the fabric configuration.

If the workload is already balanced and the I/O Busy Percentage remains high, consider installing more ports in the DS8000 system, or consider moving some of the workload to other storage systems.

ALR0701W The I/O Busy Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated

the defined warning-stress boundary value of *boundary_value*.

Explanation

I/O Busy Percentage measures the percentage of I/O operations for which the port returned a 'SCSI Queue Full' or 'Busy' status to the host server. These statuses can occur for ports on DS8000 storage systems when the number of open exchanges for the port has reached the maximum value of 2000. This situation can happen when there is too much concurrent I/O being driven to the port. A threshold is defined on this metric, which can be used to compare the value of the I/O Busy Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an I/O Busy Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary defined for the related threshold.

Action

The threshold violation indicates that the physical capability of the port on the DS8000 system is being exceeded. When a 'Busy' status is returned by the port, this usually means that the requesting host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

If the I/O Busy Percentage remains high for extended periods of time, check the workload of the affected DS8000 system to ensure that it is generally balanced across all the HA ports. If the workload is not balanced across all the ports, try to balance it. Ensure that multipath device drivers are installed and functioning properly on the host servers, so that I/O activity can be spread across multiple DS8000 ports.

If multipath drivers are functioning normally, consider configuring certain ports to allow access from only certain host servers, or consider making zoning changes between host and storage system ports. Sometimes, you might need to change the cabling or the fabric configuration.

If the workload is already balanced and the I/O Busy Percentage remains high, consider installing more ports in the DS8000 system, or consider moving some of the workload to other storage systems.

ALR0702W The I/O Busy Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined warning-idle boundary value of *boundary_value*.

Explanation

I/O Busy Percentage measures the percentage of I/O operations for which the port returned a 'SCSI Queue Full' or 'Busy' status to the host server. These statuses can occur for ports on DS8000 storage systems when the number of open exchanges for the port has reached the maximum value of 2000. This situation can happen when there is too much concurrent I/O being driven to the port. A threshold is defined on this metric, which can be used to compare the value of the I/O Busy Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an I/O Busy Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary defined for the related threshold.

Action

A zero or low I/O Busy Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0703E The I/O Busy Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined critical-idle boundary value of *boundary_value*.

Explanation

I/O Busy Percentage measures the percentage of I/O operations for which the port returned a 'SCSI Queue Full' or 'Busy' status to the host server. These statuses can occur for ports on DS8000 storage systems when the number of open exchanges for the port has reached the maximum value of 2000. This situation can happen when there is too much concurrent I/O being driven to the port. A threshold is defined on this metric, which can be used to compare the value of the I/O Busy Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an I/O Busy Percentage value that is less than or equal to the critical-idle boundary defined for the related threshold.

Action

A zero or low I/O Busy Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0704E The I/O Overrun Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined critical-stress boundary value of *boundary_value*.

Explanation

I/O Overrun Percentage measures the percentage of I/O operations that had to be discarded by the port because the number of open exchanges exceeded the maximum. This can occur for ports on DS8000 storage systems when the number of open exchanges for the port has reached the maximum value of 2000. This situation can happen when there is too much concurrent I/O being driven to the port. A threshold is defined on this metric, which can be used to compare the value of the I/O Overrun Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an I/O Overrun Percentage value that is greater than equal to the critical-stress boundary defined for the related threshold.

Action

The threshold violation indicates that the physical capability of the port on the DS8000 system is being exceeded. When an I/O operation is discarded, this usually means that the requesting host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.

Check the workload of the affected DS8000 system to ensure that it is generally balanced across all the HA ports. If the workload is not balanced across all the ports, try to balance it. Ensure that multipath device drivers are installed and functioning properly on the host servers, so that I/O activity can be spread across multiple DS8000 ports.

If multipath drivers are functioning normally, consider configuring certain ports to allow access from only certain host servers, or consider making zoning changes between host and storage system ports. Sometimes, you might need to change the cabling or the fabric configuration.

If the workload is already balanced and the I/O Overrun Percentage remains high, consider installing more ports in the DS8000 system, or consider moving some of the workload to other storage systems.

ALR0705W The I/O Overrun Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined warning-stress boundary value of *boundary_value*.

Explanation

I/O Overrun Percentage measures the percentage of I/O operations that had to be discarded by the port because the number of open exchanges exceeded the maximum. This can occur for ports on DS8000 storage systems when the number of open exchanges for the port has reached the maximum value of 2000. This situation can happen when there is too much concurrent I/O being driven to the port. A threshold is defined on this metric, which can be used to compare the value of the I/O Overrun Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an I/O Overrun Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary defined for the related threshold.

Action

The threshold violation indicates that the physical capability of the port on the DS8000 system is being exceeded. When an I/O operation is discarded, this usually means that the requesting host server needs to redrive the I/O, which negatively affects the performance of the I/O for the host.

Check the workload of the affected DS8000 system to ensure that it is generally balanced across all the HA ports. If the workload is not balanced across all the ports, try to balance it. Ensure that multipath device drivers are installed and functioning properly on the host servers, so that I/O activity can be spread across multiple DS8000 ports.

If multipath drivers are functioning normally, consider configuring certain ports to allow access from only certain host servers, or consider making zoning changes between host and storage system ports. Sometimes, you might need to change the cabling or the fabric configuration.

If the workload is already balanced and the I/O Overrun Percentage remains high, consider installing more ports in the DS8000 system, or consider moving some of the workload to other storage systems.

ALR0706W The I/O Overrun Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined warning-idle boundary value of *boundary_value*.

Explanation

I/O Overrun Percentage measures the percentage of I/O operations that had to be discarded by the port because the number of open exchanges exceeded the maximum. This can occur for ports on DS8000 storage systems when the number of open exchanges for the port has reached the maximum value of 2000. This situation can happen

when there is too much concurrent I/O being driven to the port. A threshold is defined on this metric, which can be used to compare the value of the I/O Overrun Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an I/O Overrun Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary defined for the related threshold.

Action

A zero or low I/O Overrun Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0707E The I/O Overrun Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined critical-idle boundary value of *boundary_value*.

Explanation

I/O Overrun Percentage measures the percentage of I/O operations that had to be discarded by the port because the number of open exchanges exceeded the maximum. This can occur for ports on DS8000 storage systems when the number of open exchanges for the port has reached the maximum value of 2000. This situation can happen when there is too much concurrent I/O being driven to the port. A threshold is defined on this metric, which can be used to compare the value of the I/O Overrun Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has an I/O Overrun Percentage value that is less than or equal to the critical-idle boundary defined for the related threshold.

Action

A zero or low I/O Overrun Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0708E The Zero Send Buffer Credit Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined critical-stress boundary value of *boundary_value*.

Explanation

Zero Send Buffer Credit Percentage measures the percentage of time during which the port on a DS8000 storage system had depleted its send buffer credits. That is, the percentage of time when the receiving port had no credit to provide to the port on the DS8000 system. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port that is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the overall throughput of the port, and can cascade from port to port, causing congestion in the fabric. A threshold is defined on this metric, which can be used to compare the value of Zero Send Buffer Credit Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has a Zero Send Buffer Credit Percentage value that is greater than or equal to the critical-stress boundary defined for the related threshold.

The threshold violation may indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can generally be ignored, but if the percentage remains consistently high, the errors may be a cause for concern.

Action

Contact your resource vendor for troubleshooting guidelines for the storage system, and for the connected entity (storage system, switch, or server HBA). One possible cause is a mismatch in the capability (maximum speed) of host server ports and the rest of the ports in the fabric. A port that is slow or behaving unpredictably can cause buffer credit exhaustion for the port it is connected to and all other ports in its path to the storage system.

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

ALR0709W The Zero Send Buffer Credit Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined warning-stress boundary value of *boundary_value*.

Explanation

Zero Send Buffer Credit Percentage measures the percentage of time during which the port on a DS8000 storage system had depleted its send buffer credits. That is, the percentage of time when the receiving port had no credit to provide to the port on the DS8000 system. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port that is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the overall throughput of the port, and can cascade from port to port, causing congestion in the fabric. A threshold is defined on this metric, which can be used to compare the value of Zero Send Buffer Credit Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has a Zero Send Buffer Credit Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary defined for the related threshold.

The threshold violation may indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can generally be ignored, but if the percentage remains consistently high, the errors may be a cause for concern.

Action

Contact your resource vendor for troubleshooting guidelines for the storage system, and for the connected entity (storage system, switch, or server HBA). One possible cause is a mismatch in the capability (maximum speed) of host server ports and the rest of the ports in the fabric. A port that is slow or behaving unpredictably can cause buffer credit exhaustion for the port it is connected to and all other ports in its path to the storage system.

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

ALR0710W The Zero Send Buffer Credit Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined warning-idle boundary value of *boundary_value*.

Explanation

Zero Send Buffer Credit Percentage measures the percentage of time during which the port on a DS8000 storage system had depleted its send buffer credits. That is, the percentage of time when the receiving port had no credit to provide to the port on the DS8000 system. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port that is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the overall throughput of the port, and can cascade from port to port, causing congestion in the fabric. A threshold is defined on this metric, which can be used to compare the value of Zero Send Buffer Credit Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has a Zero Send Buffer Credit Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary defined for the related threshold.

Action

A zero or low Zero Send Buffer Credit Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0711E The Zero Send Buffer Credit Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined critical-idle boundary value of *boundary_value*.

Explanation

Zero Send Buffer Credit Percentage measures the percentage of time during which the port on a DS8000 storage system had depleted its send buffer credits. That is, the percentage of time when the receiving port had no credit to provide to the port on the DS8000 system. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port that is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the overall throughput of the port, and can cascade from port to port, causing congestion in the fabric. A threshold is defined on this metric, which can be used to compare the value of Zero Send Buffer Credit Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has a Zero Send Buffer Credit Percentage value that is less than or equal to the critical-idle boundary defined for the related threshold.

Action

A zero or low Zero Send Buffer Credit Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0712E The Zero Receive Buffer Credit Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined critical-stress boundary value of *boundary_value*.

Explanation

Zero Receive Buffer Credit Percentage measures the percentage of time during which the port on a DS8000 storage system had depleted its receive buffer credits. That is, the percentage of time when the port on the DS8000 system had no credit to provide to the sending port. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port that is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the overall throughput of the port, and can cascade from port to port, causing congestion in the fabric. A threshold is defined on this metric, which can be used to compare the value of Zero Receive Buffer Credit Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has a Zero Receive Buffer Credit Percentage value that is greater than or equal to the critical-stress boundary defined for the related threshold.

The threshold violation may indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can generally be ignored, but if the percentage remains consistently high, the errors may be a cause for concern.

Action

Contact your resource vendor for troubleshooting guidelines for the storage system, and for the connected entity (storage system, switch, or server HBA). One possible cause is a mismatch in the capability (maximum speed) of the storage system port and the rest of the ports in the fabric. A slow or misbehaving port can cause buffer credit exhaustion for the port it is connected to and all other ports in its path to the host server.

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

ALR0713W The Zero Receive Buffer Credit Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined warning-stress boundary value of *boundary_value*.

Explanation

Zero Receive Buffer Credit Percentage measures the percentage of time during which the port on a DS8000 storage system had depleted its receive buffer credits. That is, the percentage of time when the port on the DS8000 system had no credit to provide to the sending port. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port that is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the overall throughput of the port, and can cascade from port to port, causing congestion in the fabric. A threshold is defined on this metric, which can be used to compare the value of Zero Receive Buffer Credit Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has a Zero Receive Buffer Credit Percentage value that is greater than or equal to the warning-stress boundary, but less than the critical-stress boundary defined for the related threshold.

The threshold violation may indicate that a hardware or configuration problem exists either for this port or for the connected fabric. Isolated errors can generally be ignored, but if the percentage remains consistently high, the errors may be a cause for concern.

Action

Contact your resource vendor for troubleshooting guidelines for the storage system, and for the connected entity (storage system, switch, or server HBA). One possible cause is a mismatch in the capability (maximum speed) of the storage system port and the rest of the ports in the fabric. A slow or misbehaving port can cause buffer credit exhaustion for the port it is connected to and all other ports in its path to the host server.

Because isolated temporary errors can generally be ignored, consider enabling the alert suppression option "Suppress alerts unless the triggering condition has been violated continuously for a specified length of time" for this threshold. This can reduce the number of unnecessary alerts.

ALR0714W The Zero Receive Buffer Credit Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined warning-idle boundary value of *boundary_value*.

Explanation

Zero Receive Buffer Credit Percentage measures the percentage of time during which the port on a DS8000 storage system had depleted its receive buffer credits. That is, the percentage of time when the port on the DS8000 system had no credit to provide to the sending port. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port that is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the overall throughput of the port, and can cascade from port to port, causing congestion in the fabric. A threshold is defined on this metric, which can be used to compare the value of Zero Receive Buffer Credit Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has a Zero Receive Buffer Credit Percentage value that is less than or equal to the warning-idle boundary, but greater than the critical-idle boundary defined for the related threshold.

Action

A zero or low Zero Receive Buffer Credit Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0715E The Zero Receive Buffer Credit Percentage of port *port_name* in storage system *device_name* was found to be *measured_value*, which violated the defined critical-idle boundary value of *boundary_value*.

Explanation

Zero Receive Buffer Credit Percentage measures the percentage of time during which the port on a DS8000 storage system had depleted its receive buffer credits. That is, the percentage of time when the port on the DS8000 system had no credit to provide to the sending port. Buffer to Buffer Credit is used as a mechanism to enforce flow-control by the fibre channel protocol, and ensures that a transmitting port that is sending data does not overrun the ability of the receiving port to receive that data. The lack of buffer credits has a negative effect on the overall throughput of the port, and can cascade from port to port, causing congestion in the fabric. A threshold is defined on this metric, which can be used to compare the value of Zero Receive Buffer Credit Percentage with the defined boundaries for each set of performance statistics collected from the storage system.

A threshold boundary violation occurred. The specified storage system port has a Zero Receive Buffer Credit Percentage value that is less than or equal to the critical-idle boundary defined for the related threshold.

Action

A zero or low Zero Receive Buffer Credit Percentage value is the ideal situation. Disable the idle boundary checking by leaving the idle boundary values blank in the threshold definition for this port.

ALR0716E The number of bytes received on each of the client network interfaces of cluster *cluster_name* in storage subsystem *storage system name* was measured to be *measured_valuebytes*, which violated the defined critical-idle boundary value of *boundary_valuebytes*.

Explanation

Action

ALR1114I New Storage Subsystem discovered.

Explanation

A new storage subsystem has been discovered.

Action

No action is required.

ALR4000I Endpoint device *endpoint device name* has been discovered.

Explanation

The specified endpoint device has been discovered.

Action

No action is required.

ALR4001W Endpoint device *endpoint device name* is missing.

Explanation

The specified endpoint device has gone missing.

Action

If this alert was unexpected, check the device and it's connection to the fabric.

ALR4002I Endpoint device *endpoint device name* has been rediscovered.

Explanation

The specified endpoint device has been rediscovered.

Action

No action is required.

ALR4015I *subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN* has been discovered.

Explanation

The specified port of the entity has been discovered.

Action

No action is required.

ALR4016W *subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN* is missing.

Explanation

The specified port of the entity has gone missing.

Action

If this alert was unexpected, check the device port and it's connection to the fabric.

ALR4017I *subsystem/fabric/switch/server subsystem/fabric/switch/server name port port name or WWPN* has been rediscovered.

Explanation

The specified port of the entity has been rediscovered.

Action

No action is required.

ALR4018W *subsystem/fabric/switch/server
subsystem/fabric/switch/server name port port name or WWPN* has
gone offline.

Explanation

The specified port of the entity has gone offline.

Action

If this alert was unexpected, check the device port and it's connection to the fabric.

ALR4019I *subsystem/fabric/switch/server
subsystem/fabric/switch/server name port port name or WWPN* has
gone online.

Explanation

The specified port of the entity has gone online.

Action

No action is required.

ALR4020I *Switch switch name or WWN* has been discovered.

Explanation

The specified switch has been discovered.

Action

No action is required.

ALR4021E *Switch switch name or WWN* is missing.

Explanation

The specified switch has gone missing

Action

If this alert was unexpected, check the switch and it's connection to the fabric.

ALR4022I *Switch switch name or WWN* has been rediscovered.

Explanation

The specified switch has been rediscovered.

Action

No action is required.

ALR4023W The version for switch *switch name* or *WWN* has changed from *previous firmware version* to *current firmware version* .

Explanation

The version for the specified switch has changed .

Action

No action is required.

ALR4024W Status of switch *switch name* or *WWN* has degraded from *previous status* to *current status* .

Explanation

The status for the specified switch has degraded. This may occur, for instance, if a fan has failed or if a redundant power supply failed.

Action

Check the switch status through the switch's element manager.

ALR4025I Status of switch *switch name* or *WWN* has improved from *previous status* to *current status* .

Explanation

The status for the specified switch has improved .

Action

No action is required.

ALR4026I Blade *blade name* or *WWN* on switch *switch name* or *WWN* has been discovered.

Explanation

The specified blade on the specified switch has been discovered.

Action

No action is required.

ALR4027W Blade *blade name* or *WWN* on switch *switch name* or *WWN* is missing.

Explanation

The specified blade on the specified switch has gone missing.

Action

If this alert was unexpected, check the blade on the switch.

ALR4028I Blade *blade name or WWN* on switch *switch name or WWN* has been rediscovered.

Explanation

The specified blade on the specified switch has been rediscovered.

Action

No action is required.

ALR4029E Blade *blade name or WWN* on switch *switch name or WWN* has gone offline.

Explanation

The specified blade on the specified switch has gone offline.

Action

If this alert was unexpected, check the blade on the switch.

ALR4030I Blade *blade name or WWN* on switch *switch name or WWN* has gone online.

Explanation

The specified blade on the specified switch has gone online.

Action

No action is required.

ALR4034W The driver version for HBA *adapter name* on server *host name* has changed from *previous version* to *new version*.

Explanation

IBM Spectrum Control detected that the driver version for an HBA on a monitored server or server has changed.

Action

This message is for informational purposes only. No further action is required.

ALR4035W The firmware version for HBA *adapter name* on server *host name* has changed from *previous version* to *new version* .

Explanation

IBM Spectrum Control detected that the firmware version for an HBA on a monitored server or server has changed.

Action

This message is for informational purposes only. No further action is required.

ALR4046I Fabric *fabric name or WWN* is discovered.

Explanation

The specified fabric has been discovered.

Action

No action is required.

ALR4047E Fabric *fabric name or WWN* is missing.

Explanation

The specified fabric has gone missing. This alert may indicate that the principal switch on the fabric has changed, or that the Data Sources managing the fabric have lost communication with the fabric.

Action

If this alert was unexpected, check if another fabric with the same set of switches has been discovered. This would indicate that the principal switch changed on the fabric. Otherwise, check the status of the Data Sources which were managing the fabric.

ALR4048I Fabric *fabric name or WWN* is rediscovered.

Explanation

The specified fabric has been rediscovered.

Action

No action is required.

ALR4051I Inactive zone *zone name* in fabric *fabric name or WWN* has been discovered.

Explanation

The specified inactive zone in the specified fabric has been discovered.

Action

No action is required.

ALR4052W Inactive zone *zone name* in fabric *fabric name or WWN* is missing.

Explanation

The specified inactive zone in the specified fabric has gone missing.

Action

No action is required.

ALR4053I Inactive zone *zone name* in fabric *fabric name or WWN* has been rediscovered.

Explanation

The specified inactive zone in the specified fabric has been rediscovered.

Action

No action is required.

ALR4054I Inactive zoneset *zoneset name* in fabric *fabric name* or *WWN* has been discovered.

Explanation

The specified inactive zoneset in the specified fabric has been discovered.

Action

No action is required.

ALR4055W Inactive zoneset *zoneset name* in fabric *fabric name* or *WWN* is missing.

Explanation

The specified inactive zoneset in the specified fabric has gone missing.

Action

No action is required.

ALR4056I Inactive zoneset *zoneset name* in fabric *fabric name* or *WWN* has been rediscovered.

Explanation

The specified inactive zoneset in the specified fabric has been rediscovered.

Action

No action is required.

ALR4063I The connection from *switch or node switch name or WWN port port name or WWPN* to *switch or node switch name or WWN port port name or WWPN* has been discovered.

Explanation

The connection from the specified port of the specified switch or node to the specified port of the specified switch or node has been discovered.

Action

No action is required.

ALR4064W The connection from *switch or node switch name or WWN port port name or WWPN* to *switch or node switch name or WWN port port name or WWPN* is missing.

Explanation

The connection from the specified port of the specified switch or node to the specified port of the specified switch or node has gone missing.

Action

If this alert was unexpected, check the connection between the ports and check the status of the devices at both ends of the connection.

ALR4065I The connection from *switch or node switch name or WWN port port name or WWPN* to *switch or node switch name or WWN port port name or WWPN* has been rediscovered.

Explanation

The connection from the specified port of the specified switch or node to the specified port of the specified switch or node has been rediscovered.

Action

No action is required.

ALR4066I Switch *switch name or WWN* in fabric *fabric name or WWN* has been discovered.

Explanation

The specified switch in the specified fabric has been discovered.

Action

No action is required.

ALR4067W Switch *switch name or WWN* in fabric *fabric name or WWN* is missing.

Explanation

The specified switch in the specified fabric has gone missing.

Action

If this alert was unexpected, check the status of the switch and it's connections to the other switches in the fabric.

ALR4068I Switch *switch name or WWN* in fabric *fabric name or WWN* has been rediscovered.

Explanation

The specified switch in the specified fabric has been rediscovered.

Action

No action is required.

ALR4069I Port *port name or WWPN* in switch *switch name or WWN* has been discovered.

Explanation

The specified port in the specified switch has been discovered.

Action

No action is required.

ALR4070W Port *port name or WWPN* in switch *switch name or WWN* is missing.

Explanation

The specified port in the specified switch has gone missing.

Action

If this alert was unexpected, check the status of the switch. If the switch is a logical switch in a Virtual Fabric, check the port assignments for the Virtual Fabric.

ALR4071I Port *port name or WWPN* in switch *switch name or WWN* has been rediscovered.

Explanation

The specified port in the specified switch has been rediscovered.

Action

No action is required.

ALR4078I Alias *zone alias* has been added to inactive zone *zone name* in fabric *fabric name or WWN* .

Explanation

The specified alias has been added to the specified inactive zone in the specified fabric.

Action

No action is required.

ALR4079W Alias *zone alias* has been removed from inactive zone *zone name* in fabric *fabric name or WWN* .

Explanation

The specified alias has been removed from the specified inactive zone in the specified fabric.

Action

No action is required.

ALR4080I Alias *zone alias* has been readded to inactive zone *zone name* in fabric *fabric name or WWN* .

Explanation

The specified alias has been readded to the specified inactive zone in the specified fabric.

Action

No action is required.

ALR4081I Zone member *zone member name* has been added to inactive zone *zone name* in fabric *fabric name* or *WWN* .

Explanation

The specified zone member has been added to the specified inactive zone in the specified fabric.

Action

No action is required.

ALR4082I Zone member *zone member name* has been removed from inactive zone *zone name* in fabric *fabric name* or *WWN* .

Explanation

The specified zone member has been removed from the specified inactive zone in the specified fabric.

Action

No action is required.

ALR4083I Zone member *zone member name* has been readded to inactive zone *zone name* in fabric *fabric name* or *WWN* .

Explanation

The specified zone member has been readded to the specified inactive zone in the specified fabric.

Action

No action is required.

ALR4084I Zone *zone name* has been added to inactive zone set *zone set name* in fabric *fabric name* or *WWN* .

Explanation

The specified zone has been added to the specified inactive zone set in the specified fabric.

Action

No action is required.

ALR4085I Zone *zone name* has been removed from inactive zone set *zone set name* in fabric *fabric name* or *WWN* .

Explanation

The specified zone has been removed from the specified inactive zone set in the specified fabric.

Action

No action is required.

ALR4086I Zone *zone name* has been readdded to inactive zone set *zone set name* in fabric *fabric name* or *WWN* .

Explanation

The specified zone has been readdded to the specified inactive zone set in the specified fabric.

Action

No action is required.

ALR4089W ZoneSet *zoneset name* in fabric *fabric name* or *WWN* has been deactivated. ZoneSet *zoneset name* has been activated.

Explanation

The specified zoneset in the specified fabric has been deactivated. Another zoneset has been activated.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4090W Active zone *zone name* in fabric *fabric name* or *WWN* is missing.

Explanation

The specified active zone in the specified fabric has gone missing.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4091I Active zone *zone name* in fabric *fabric name* or *WWN* has been discovered.

Explanation

The specified active zone in the specified fabric has been discovered.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4092I Active zoneset *zoneset name* in fabric *fabric name* or *WWN* has been discovered.

Explanation

The specified active zoneset in the specified fabric has been discovered.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4093E ZoneSet *zoneset name* in fabric *fabric name* or *WWN* has been deactivated. ZoneSet *zoneset name* has been activated.

Explanation

The specified zoneset in the specified fabric has been deactivated. No zoneset has been activated..

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4094I Active zone *zone name* in fabric *fabric name* or *WWN* has been rediscovered.

Explanation

The specified active zone in the specified fabric has been rediscovered.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4095I Active zoneset *zoneset name* in fabric *fabric name* or *WWN* has been rediscovered.

Explanation

The specified active zoneset in the specified fabric has been rediscovered.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4096I Zone member *zone member name* has been added to active zone *zone name* in fabric *fabric name* or *WWN* .

Explanation

The specified zone member has been added to the specified active zone in the specified fabric.

Action

No action is required.

ALR4097I Zone member *zone member name* has been removed from active zone *zone name* in fabric *fabric name* or *WWN* .

Explanation

The specified zone member has been removed from the specified active zone in the specified fabric.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4098I Zone member *zone member name* has been readded to active zone *zone name* in fabric *fabric name* or *WWN* .

Explanation

The specified zone member has been readded to the specified active zone in the specified fabric.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4099I Zone *zone name* has been added to active zone set *zone set name* in fabric *fabric name* or *WWN* .

Explanation

The specified zone has been added to the specified active zone set in the specified fabric.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4103W The performance monitor's primary process is experiencing a high memory utilization.

Explanation

The IBM Spectrum Control server encountered a high memory utilization for a period of time, which may lead to stability problems.

Action

Check, if there was high activity on the server. If so, try to reduce the workload, for example by rescheduling probes or batch reports.

ALR4104W A database used by the system is reporting an alarm: *value*.

Explanation

The system database or the database manager hosting the system repository is reporting an alarm. The message returns a value of:

- DBM - database manager
- DB - repository database
- DBM, DB - both

Action

Check the status of the affected database.

ALR4105W Device server configuration should be changed to improve performance: *value*.

Explanation

A device server workload queue is reporting an alert. The value returned with this message identifies the queue. The queue length can be reduced by adjusting the device server configuration.

Action

Send this message, including the value, to IBM Technical Support.

Related reference

-  [Getting support](#)

ALR4106W The IBM Spectrum Control server is receiving a high number of external type of events received, which is either CIM for CIM indications or SNMP for SNMP traps. events, which may cause temporary performance degradation.

Explanation

The IBM Spectrum Control server encountered a high rate of incoming external events, which are potentially affecting IBM Spectrum Control performance.

Action

Check the environment for items, which are producing a unusual high number of events and correct the problem.

ALR4107I Zone *zone name* has been removed from active zone set *zone set name* in fabric *fabric name* or *WWN* .

Explanation

The specified zone has been removed from the specified active zone set in the specified fabric.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4108I The server *server name* at host *host name* successfully connected to the database after previous attempts failed.

Explanation

The specified server successfully connected to the database after previous attempts failed.

Action

No action is required.

ALR4112E The server *server name* at host *host name* failed to connect to the database.

Explanation

The database is not available. Restart the database if it is offline.

Action

Check that the database is operational and online, else contact IBM customer technical support.

Related reference

-  [Getting support](#)

ALR4109I Alias *zone alias* has been added to active zone *zone name* in fabric *fabric name* or *WWN* .

Explanation

The specified alias has been added to the specified active zone in the specified fabric.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4113E The IBM Spectrum Control server failed to connect to the database .

Explanation

An error occurred when IBM Spectrum Control tried to connect to the repository database.

Action

Verify that the database is operational and online.

ALR4110W Alias *alias name* has been removed from active zone *active zone name* in fabric *fabric name* .

Explanation

The specified alias is not available, because it has been removed from the active zone.

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4111I Alias *alias name* has been readdded to active zone *active zone name* in fabric *fabric name* .

Explanation

The specified alias has been readdded to the specified active zone in the specified fabric

Action

If this alert was unexpected, check the active zoneset definition of your fabric.

ALR4197W A new connection is detected.

Explanation

A new fabric connection is detected.

Action

No action is required.

ALR4198W The state for connection from *initial state* to *final state* has changed.

Explanation

A fabric connection is missing or rediscovered.

Action

If this alert was unexpected, check the connection between devices in the fabric.

ALR4199W The state for Fabric *fabric name* has changed.

Explanation

The specified fabric is missing or rediscovered.

Action

If this alert was unexpected, check if another fabric with the same set of switches has been discovered or gone missing. This would indicate that the principal switch changed on the fabric. Otherwise, check the status of the Data Sources which were managing the fabric.

ALR4224W The state has changed for Node *node name* .

Explanation

The state has changed for the specified node.

Action

No action is required.

ALR4225I Node *node name* has been discovered.

Explanation

The specified node has been discovered.

Action

No action is required.

ALR4226W Node *node name* has gone offline.

Explanation

The specified node is offline.

Action

No action is required.

ALR4227I Node *node name* has gone online.

Explanation

The specified node is online.

Action

No action is required.

ALR4241E Subsystem *storage subsystem name* has gone offline.

Explanation

The specified storage subsystem is offline.

Action

No action is required.

ALR4242I Subsystem *storage subsystem name* has gone online.

Explanation

The specified storage subsystem is online.

Action

No action is required.

ALR4243W The subsystem version has changed from *initial version* to *new version* on Subsystem *storage subsystem name* .

Explanation

The specified storage subsystem version has changed.

Action

No action is required.

ALR4244W The allocated capacity has changed from *initial capacity* to *new capacity* on Subsystem *storage subsystem name* .

Explanation

The specified storage subsystem allocated capacity has changed.

Action

No action is required.

ALR4245W The available capacity has changed from *initial capacity* to *new capacity* on Subsystem *storage subsystem name* .

Explanation

The specified storage subsystem available capacity has changed.

Action

No action is required.

ALR4246W Back-end capacity has changed from *initial capacity* to *new capacity* on Subsystem *storage subsystem name* .

Explanation

The specified storage subsystem back-end capacity has changed.

Action

No action is required.

ALR4247W Back-end controller *back-end controller name* for owning storage *subsystem name* has gone offline.

Explanation

The specified back-end controller on the specified subsystem is offline.

Action

No action is required.

ALR4248I Back-end controller *back-end controller name* for owning storage *subsystem name* has gone online.

Explanation

The specified back-end controller on the specified subsystem is online.

Action

No action is required.

ALR4249W Volume *volume name* on storage *subsystem name* has gone offline.

Explanation

The specified volume on the specified subsystem is offline.

Action

No action is required.

ALR4250I Volume *volume name* on storage *subsystem name* has gone online.

Explanation

The specified volume on the specified subsystem is online.

Action

No action is required.

ALR4251W The capacity has changed from *initial capacity* to *new capacity* for Volume *volume name* on Subsystem storage *subsystem name* .

Explanation

The capacity has changed for the specified volume on the specified subsystem .

Action

No action is required.

ALR4252W The state for Pool *pool name* on Subsystem storage *subsystem name* has changed to not detectable .

Explanation

The state for the specified pool on the specified subsystem has changed to not detectable.

Action

No action is required.

ALR4253I Pool *pool name* on Subsystem storage *subsystem name* has been discovered.

Explanation

The specified pool on the specified subsystem has been discovered.

Action

No action is required.

ALR4254W Pool *pool name* on Subsystem storage *subsystem name* has gone offline.

Explanation

The specified pool on the specified subsystem is offline.

Action

No action is required.

ALR4255I Pool *pool name* on Subsystem storage *subsystem name* has gone online.

Explanation

The specified pool on the specified subsystem is online.

Action

No action is required.

ALR4256W The pool capacity has changed from *initial capacity* to *new capacity* for Pool *pool name* on Subsystem storage *subsystem name* .

Explanation

The pool capacity has changed for the specified pool on the specified subsystem .

Action

No action is required.

ALR4257W The pool available space has changed from *initial capacity* to *new capacity* for Pool *pool name* on Subsystem storage *subsystem name* .

Explanation

The pool available space has changed for the specified pool on the specified subsystem .

Action

No action is required.

ALR4273E Server *server name* has gone offline.

Explanation

The specified server has gone offline .

Action

If this alert was unexpected, check the specified server.

ALR4274I Server *server name* has gone online.

Explanation

The specified server has gone online .

Action

No action is required.

ALR4278W The property for Subsystem storage *subsystem name* has changed.

Explanation

A port on the specified subsystem is newly discovered, missing, or rediscovered.

Action

No action is required.

ALR4300W The use count for Disk Drive *disk drive name* on Subsystem *subsystem name* has changed from *initial use count* to *final use count*.

Explanation

The specified disk drive's use count has changed as mentioned in the message description.

Action

No action is required.

ALR4301W Disk Drive *disk drive name* on Subsystem *subsystem name* has gone offline.

Explanation

The specified disk drive on the specified subsystem is offline.

Action

No action is required.

ALR4304W The state for Back-end Controller *back-end controller* for *subsystem* has changed from *initial state* to *final state* .

Explanation

The specified back-end controller is missing or has been rediscovered.

Action

No action is required.

ALR4305W The WWPN path count for Back-end Controller *back-end controller* for *system name* has changed from *initial state* to *final state* .

Explanation

The WWN path count for the specified Back-end Controller has changed.

Action

No action is required.

ALR4312W Notification has received from external device *device name*

Explanation

A notification was received from the specified external device.

Action

No action is required.

ALR4313W The endpoint version has changed from *initial state* to *final state* on Endpoint

Explanation

The specified endpoint version has been changed.

Action

No action is required.

ALR4314W Entity *entity name* has been discovered.

Explanation

A new entity (specified in the message) has been discovered.

Action

No action is required.

ALR4315I Hypervisor *hypervisor name* has been discovered.

Explanation

A new hypervisor has been discovered.

Action

No action is required.

ALR4316I Virtual machine *vm name* was added to hypervisor *hypervisor name* .

Explanation

The specified virtual machine has been added.

Action

No action is required.

ALR4317W Virtual machine *vm name* was removed from hypervisor *hypervisor name* .

Explanation

The specified virtual machine has been removed.

Action

No action is required.

ALR4318W New unmanaged hypervisor discovered.

Explanation

A new hypervisor has been discovered, it is not yet managed.

Action

No action is required.

ALR4319W Virtual machine added.

Explanation

A virtual machine was added.

Action

No action is required.

ALR4320W Virtual machine removed.

Explanation

A virtual machine was removed.

Action

No action is required.

ALR4321W Hypervisor *hypervisor* missing.

Explanation

The specified hypervisor is missing.

Action

No action is required.

ALR4322W Hypervisor missing.

Explanation

A hypervisor is missing.

Action

No action is required.

ALR4323I New disk *disk name* discovered for system *system name* .

Explanation

A new disk has been discovered.

Action

No action is required.

ALR4324W Disk *disk* not found for system *system* .

Explanation

The specified disk was not found for the mentioned system.

Action

No action is required. If the disk is expected to be found, please contact IBM customer technical support.

Related reference

- [Getting support](#)

ALR4325W New volume *volume name* discovered for system *system name*

.

Explanation

A new volume has been discovered on the specified system.

Action

No action is required.

ALR4326W Volume *volume* not found for system *system*.

Explanation

The specified volume is not found on the system.

Action

No action is required. If the specified volume is expected to be found, please contact IBM customer technical support.

Related reference

- [Getting support](#)

ALR4327W Zone Alias to Member Change

Explanation

A zone alias to member association is discovered, rediscovered, or is missing.

Action

No action is required.

ALR4328W The association between Zone Alias *zone alias* and Member *member* has changed.

Explanation

The specified zone alias does not include the specified member anymore.

Action

No action is required.

ALR4329I Zone member *zone member* has been added to zone alias *zone alias* in fabric *fabric* .

Explanation

The specified zone member has been added to the specified zone alias in the fabric.

Action

No action is required.

ALR4330W Zone member *zone member* has been removed from zone alias *zone alias* in fabric *fabric* .

Explanation

The specified zone member has been removed from the specified zone alias in the fabric.

Action

No action is required.

ALR4331I Zone member *zone member* has been readded to zone alias *zone alias* in fabric *fabric* .

Explanation

The specified zone member has been readded to the specified zone alias in the fabric.

Action

No action is required.

ALR4332W Storage Resource Agent Deployment Failed

Explanation

The deployment of the Storage Resource Agent has failed.

Action

Please check the logs and follow the SRA documentation related to the error code from the logs. Contact IBM customer technical support if this does not solve the problem.

Related reference

- [Getting support](#)

ALR4333W Replication Session State Change alert received.

Explanation

The replication session state has changed.

Action

No action is required.

ALR4334W Replication Configuration Change alert received.

Explanation

The replication configuration has changed.

Action

No action is required.

ALR4335W Replication Suspending Event Notification alert received.

Explanation

Replication Suspending Event Notification alert.

Action

No action is required.

ALR4336W Replication Communication Failure alert received.

Explanation

The replication communication has failed.

Action

No action is required.

ALR4337W Replication Management Server State Change alert received

Explanation

The Replication Management server state has changed.

Action

No action is required.

ALR4338W Replication PPRC Path State Change alert received.

Explanation

Replication PPRC Path state has changed.

Action

No action is required.

ALR4339W The IBM Spectrum Control for Replication resource *resource nametriggered an alert with the following message:message text*

Explanation

Tivoli Storage Productivity Center for replication resource has triggered an alert with the specified message.

For more information about replication alerts, see IBM Spectrum Control Help -> Alerting -> Triggering conditions for alerts -> Triggering conditions for replication alerts in the online help.

Action

No action is required.

ALR4353W *number of affected datapaths* Data Paths from Host *host name* to Volume *volume name* on Subsystem *subsystem name* are no longer available.

Explanation

The specified number of datapaths from the specified host to the specified volume on the specified subsystem are no longer available.

Action

No action is required.

ALR4354I *number of affected datapaths* Data Paths from Host *host name* to Volume *volume name* on Subsystem *subsystem name* have been discovered..

Explanation

The specified number of datapaths from the specified host to the specified volume on the specified subsystem have been discovered.

Action

No action is required.

ALR4356E The mount state of specified file system changed to error level.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSSA0619C. This indicates that the system has detected a significant change in the use of a monitored resource. See the information center for the device for details.

Action

Correct the error or contact the device system administrator.

ALR4358I The mount state of specified file system changed to normal level.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSSA0621I. This indicates that the system has detected a significant change in the use of a monitored resource. See the information center for the device for details.

Action

No action is required.

ALR4359E The CPU usage reached the error level.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSS1D0153C. This indicates that the CPU usage exceeded the error level. See the information center for the device for details.

Action

No action is required.

ALR4360W The CPU usage reached the warning level.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSSF1D0152W. This indicates that the CPU usage exceeded the warning level. See the information center for the device for details.

Action

No action is required.

ALR4361I The CPU usage reached the normal level.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSSF1D0151I. This indicates that the CPU usage has returned to a level below the warning threshold. See the information center for the device for details.

Action

No action is required.

ALR4362E The memory usage reached the error level.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSS200825C. This indicates that the system memory usage exceeded the error level. See the information center for the device for details.

Action

Contact the device administrator.

ALR4364I The memory usage reached the normal level.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSSF200828I. This indicates that the system memory usage has returned to a level below the warning threshold. See the information center for the device for details.

Action

No action is required.

ALR4365I The clustered CIFS is active.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSSF0010I. This indicates that the clustered CIFS is active. See the information center for the device for details.

Action

No action is required.

ALR4366W The clustered CIFS is disabled.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSSF0011W. This indicates that the clustered CIFS is disabled. See the information center for the device for details.

Action

No action is required.

ALR4367E The clustered CIFS status reached the error level.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSSF0011W. This indicates that the clustered CIFS status reached the error level. See the information center for the device for details.

Action

Contact the device system administrator.

ALR4368I The IBM Spectrum Scale is active.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSSA0046I. This indicates that the IBM Spectrum Scale is active. See the information center for the device for details.

Action

No action is required.

ALR4369W The IBM Spectrum Scale status reached the warning level.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSSA0050W. This indicates that the IBM Spectrum Scale status reached the warning level. See the information center for the device for details.

Action

No action is required.

ALR4370E The IBM Spectrum Scale is down.

Explanation

The IBM Spectrum Control server received an SNMP trap from the NAS device with the message EFSSA0047C. This indicates that the IBM Spectrum Scale is down. See the information center for the device for details.

Action

Contact the device administrator.

ALR4427I The file system *file system* has been detected on device *device display name*.

Explanation

The IBM Spectrum Control server found a filesystem on the device for the first time.

Action

No action required.

ALR4429I The capacity of file system *file system* has changed from *previous capacity* to *current capacity* on device *device NAS display name*.

Explanation

The capacity of the file system has changed.

Action

No action required.

ALR4433W The free space on file system *file system* has fallen below the threshold value of *threshold* on device *device nas display name*. The free space is *current free space value* or *relative free space value* of the file system capacity.

Explanation

A user defined file system threshold has been exceeded.

Action

No action required.

ALR4440W The state of the device type *node type node node name* changed from *old state* to *new state* on device *device display name*.

Explanation

The state of a node has changed.

Action

Contact the device administrator.

ALR4441I The state of the device type *node type node node name* changed from *old state* to *new state* on device *device display name*.

Explanation

The state of a node has changed.

Action

No action required.

ALR4442W The IBM Spectrum Scale status of *device type* node *node name* changed from *old state* to *new state* on device *display name*.

Explanation

The state of a node has changed.

Action

Contact the device administrator.

ALR4443I The IBM Spectrum Scale status of *device type* node *node name* changed from *old state* to *new state* on device *display name*.

Explanation

The state of a node has changed.

Action

No action required.

ALR4447I The fileset *fileset* was detected on *device type* device *device name*.

Explanation

The IBM Spectrum Control server found a fileset for the first time.

Action

No action required.

ALR4448I The fileset *fileset* was linked to *path* for *device type* device *device name*.

Explanation

The fileset was linked.

Action

No action required.

ALR4455I The fileset *fileset* was unlinked on *device type* device *device name*.

Explanation

The fileset was unlinked.

Action

No action required.

ALR4458W The number of free inodes on file system *path* has fallen below the threshold value of *threshold* on device *type device device name*. There are *current value* free inodes or *current value relative to maximum* of the maximum inodes.

Explanation

The file system is low on inodes.

Action

No action required.

ALR4460I Export *export name* detected on device *type device device name* with path *path*.

Explanation

A new export was detected for the network attached storage device.

Action

No action required.

ALR4461W The state of export *export name* changed from *previous state* to *current state* on device *type device device name*.

Explanation

The state of the export changed.

Action

No action required.

ALR4462W Export *export name* was reconfigured on device *type device device name*. Path changed from *previous path* to *current path*.

Explanation

The export was linked to a new path.

Action

No action required.

ALR4463W Export *export name* was reconfigured on device *type device device name*. Protocols changed from *previous list of protocols* to *current list of protocols*.

Explanation

The protocols configured for the share have been changed.

Action

No action required.

ALR4470W Export *export name* is missing from *device type* *device name*..

Explanation

An export that was previously found during a probe was not found in the current probe.

Action

No action required.

ALR4471I Export *export name* was rediscovered on *device type* *device name*..

Explanation

An export that was previously missing was found again by the current probe.

Action

No action required.

ALR4474W Fileset *fileset name* is missing from *device type* *device name*..

Explanation

A fileset that was previously found during a probe was not found in the current probe.

Action

No action required.

ALR4475I Fileset *fileset name* was rediscovered on *device type* *device name*..

Explanation

A fileset that was previously missing was found again by the current probe.

Action

No action required.

ALR4478W File system *File system name* is missing from *device type* *device name*..

Explanation

A file system that was previously found during a probe was not found in the current probe.

Action

No action required.

ALR4479I File system *File system name* was rediscovered on device *device device name*..

Explanation

A file system that was previously missing was found again by the current probe.

Action

No action required.

ALR4482W A *Quota type* *Quota limit type* quota was violated for the *path* file system on the device *device name device type* system. *Quota type resource name* is consuming usage and the *Quota limit type limit* is *threshold*.

Explanation

A quota violation was detected on the NAS system. Quotas define limits on the amount of storage and inodes that a user, group of users, or fileset can consume. Violations of quotas are detected when a probe schedule collects data about the NAS system using the `lsquota` command. By default, the NAS system will update these values daily. See the NAS documentation for details.

Action

Reduce the file system use or contact the device administrator to increase the quota limits.

ALR4496I New quota detected on file system *path* of *Device type* device *Device name*.

Explanation

A quota that was defined on the NAS has been detected for the first time by the IBM Spectrum Control server.

Action

No action required.

ALR4385E The status of NSD *NSD name* reached error level.

Explanation

NSD usage has reached an error level.

Action

Contact the device administrator.

ALR4386W The status of NSD *NSD name* reached warning level.

Explanation

NSD usage has reached a warning level.

Action

Contact the device administrator.

ALR4387I The status of NSD *NSD name* was set back to normal level.

Explanation

NSD usage was set back to normal level.

Action

No action required.

ALR4503I New NSD *NSD name* has been detected on device type *device* device display name.

Explanation

The IBM Spectrum Control server found an NSD on the device for the first time.

Action

No action required.

ALR4505W NSD *NSD name* is missing from device type *device* device name.

Explanation

An NSD that was previously found during a probe was not found in the current probe.

Action

No action required.

ALR4507W The state of NSD *NSD name* changed from *previous state* to *current state* on device type *device* device name.

Explanation

The state of the NSD has changed.

Action

No action required.

ALR4511E Alert condition for nodes has been selected. Select only nodes.

Explanation

When an alert condition for nodes is selected, you need to specify nodes.

Action

Specify one of more nodes.

ALR4512E Alert condition for clusters has been selected. Select only clusters.

Explanation

When an alert condition for clusters is selected, you need to specify clusters.

Action

Specify one of more clusters.

ALR4513E Alert condition for NSD has been selected. Select only NSD.

Explanation

When an alert condition for NSDs is selected, you need to specify NSDs.

Action

Specify one of more NSDs.

ALR4514E Alert condition for File set has been selected. Select only File set.

Explanation

When an alert condition for File sets is selected, you need to specify File sets.

Action

Specify one of more File sets.

ALR1022M A new unmanaged server or cluster discovered.

Explanation

A new unmanaged server or cluster discovered.

Action

No action required.

ALR1294W The server or cluster has gone offline.

Explanation

The server or cluster has gone offline.

Action

Contact the device administrator.

ALR1295W The server or cluster has gone online.

Explanation

The server or cluster has gone online.

Action

No action required.

ALR1296W The server or cluster property has changed.

Explanation

The server or cluster property has changed.

Action

Contact the device administrator.

ALR1245W A node state has changed.

Explanation

A node state has changed.

Action

Contact the device administrator.

ALR1246W A node was discovered.

Explanation

A node was discovered.

Action

No action required.

ALR4528I Cluster was discovered.

Explanation

The IBM Spectrum Control server found a cluster for the first time.

Action

No action required.

ALR4529I Cluster was removed.

Explanation

A cluster that was previously found during a probe was not found in the current probe.

Action

Contact the device administrator.

ALR4530I Cluster was rediscovered.

Explanation

A cluster that was previously missing was found again by the current probe.

Action

No action required.

ALR0078W =Performance monitor for device *value* failed to collect new data using data source *value*.

Explanation

The performance monitor for the specified device encountered an error and could not continue collecting performance data using specified data source

Action

Check to ensure that the data source is still operational, and restart it if necessary. Check the performance monitor job log for additional information.

ALR4391I Node *node name* is selected as cache gateway node.

Explanation

A node is selected as cache gateway node.

Action

No action required.

ALR4392I Node *node name* is unselected as cache gateway node.

Explanation

A node is unselected as cache gateway node.

Action

No action required.

ALR4393I Home system *home system name* detected on device *type* device *device name* with path *path*.

Explanation

A new home system was detected for the network attached storage device.

Action

No action required.

ALR4394I Home system has been removed from fileset *fileset name* on device *type* device *device name*.

Explanation

A home system was removed for the network attached storage device.

Action

No action required.

ALR4395W Home system *home system name* is missing from device type *device name*.

Explanation

A home system that was previously found during a probe was not found in the current probe.

Action

No action required.

ALR4396I Home system *home system name* was rediscovered on device type *device name*..

Explanation

A home system that was previously missing was found again by the current probe.

Action

No action required.

ALR4397I Cache fileset *cache fileset name* detected on device type *device name*.

Explanation

A new cache fileset was detected for the network attached storage device.

Action

No action required.

ALR4398W Cache fileset *cache fileset name* is missing from device type *device name*.

Explanation

A cache fileset that was previously found during a probe was not found in the current probe.

Action

No action required.

ALR4399I Cache fileset *cache fileset name* was rediscovered on device type *device name*.

Explanation

A cache fileset that was previously missing was found again by the current probe.

Action

No action required.

ALR4400I Cache fileset name has changed from *cache fileset name* to *cache fileset name* on device type device device name.

Explanation

A cache fileset name has changed.

Action

No action required.

ALR4401I Cache fileset *cache fileset name* state has changed from *old value* to *new value* on device type device device name.

Explanation

A cache fileset state has changed.

Action

No action required.

ALR4402I Cache fileset *cache fileset name* mode has changed from *old value* to *new value* on device type device device name.

Explanation

A cache fileset mode has changed.

Action

No action required.

ALR4403I Cache client *cluster name* is added to home system *home system name* on device type device device name.

Explanation

A cache client is added to a home system.

Action

No action required.

ALR4404I Cache client *cluster name* has been removed from home system *home system name* on device type device device name.

Explanation

A cache client is removed from a home system.

Action

No action required.

ALR4541E The available space is too low for pool *pool name* on storage system *storage system name*. The measured value *pool available space* violates the critical boundary of *user defined threshold value*.

Explanation

The thin provisioned pool available space is lower than a user defined critical boundary.

Action

No action required.

ALR4542W The available space is too low for pool *pool name* on storage system *storage system name*. The measured value *pool available space* violates the warning boundary of *user defined threshold value*.

Explanation

The thin provisioned pool available space is lower than a user defined warning boundary.

Action

No action required.

ALR4543E The allocation is too high for pool *pool name* on storage system *storage system name*. The measured value *pool virtual allocation* violates the critical boundary of *user defined threshold value*.

Explanation

The thin provisioned pool virtual allocation is higher than a user defined critical boundary.

Action

No action required.

ALR4544W The allocation is too high for pool *pool name* on storage system *storage system name*. The measured value *pool virtual allocation* violates the warning boundary of *user defined threshold value*.

Explanation

The thin provisioned pool virtual allocation is higher than a user defined warning boundary.

Action

No action required.

ALR4545E The shortfall percentage is too high for pool *pool name* on storage system *storage system name*. The measured value *pool shortfall percentage* violates the critical boundary of user defined threshold value.

Explanation

The thin provisioned pool shortfall percentage is higher than a user defined critical boundary.

Action

No action required.

ALR4546W The shortfall percentage is too high for pool *pool name* on storage system *storage system name*. The measured value *pool shortfall percentage* violates the warning boundary of user defined threshold value.

Explanation

The thin provisioned pool shortfall percentage is higher than a user defined warning boundary.

Action

No action required.

ALR4547I VMWare Cluster *cluster name* discovered on hypervisor *hypervisor name*.

Explanation

A VMware cluster discovery occurs when you add a vCenter as a data source, and one or more of the discovered hypervisors are members of a cluster. If a hypervisor that was added through the vCenter has become a member of a cluster after initial discovery through the vCenter, the cluster is discovered during the next data collection.

Action

No action required.

ALR4548W VMWare Cluster *cluster name* removed from hypervisor *hypervisor name*.

Explanation

A hypervisor that was added through the vCenter was removed as a member of a cluster after initial discovery through the vCenter, and was not added to another cluster.

Action

Check the VMWare environment to ensure that it was correct to remove the hypervisor from the cluster.

ALR4549I New cluster hypervisor relationship discovered.

Explanation

A new cluster hypervisor relationship has been discovered.

Action

No action required.

ALR4550W Cluster hypervisor relationship removed.

Explanation

A cluster hypervisor relationship has been removed.

Action

No action required.

ALR4551I Hypervisor *hypervisor name* was moved from VMWare Cluster *old cluster name* to VMWare Cluster *new cluster name*.

Explanation

A hypervisor that was added through the vCenter was removed as a member of a cluster after initial discovery through the vCenter, and has been added to another cluster.

Action

No action required.

ALR4552I Cluster hypervisor relationship moved.

Explanation

A cluster hypervisor relationship has been moved.

Action

No action required.

ALR4600I Fabric Name *fabric WWN* changed to *fabric WWN* .

Explanation

The specified fabric has changed WWN, due to Principal Switch change.

Action

No action is required.

ALR1349I A new path *path name* was discovered for disk *disk name* on host *host name*.

Explanation

A new access path was discovered for the specified disk.

Action

No action required.

ALR1350W The path *path name* was not found for disk *disk name* on host *host name*.

Explanation

An access path for a disk that was discovered previously can no longer be found.

Action

No action required.

ALR1351W The path *path name* for disk *disk name* on host *host name* is disconnected.

Explanation

The path is disconnected. This change might or might not affect the availability of the disk because there might be more than one path available.

Action

Check the configuration. Ensure that the volume associated with the path is still assigned to the host. Ensure that the SAN is still zoned correctly between the storage system and the host ports.

ALR4604I The *home system* home system was linked to *path* on resource type resource *resource name*.

Explanation

The home system fileset was linked to a path.

Action

This message is for informational purposes only. No action is required.

ALR4605I The *home system* home system was unlinked from a path on resource type resource *resource name*.

Explanation

The home system fileset was unlinked from a path.

Action

This message is for informational purposes only. No action is required.

ALR1352E The status of disk *disk name* on server *server name* has degraded to Error from *former status*.

Explanation

At least one permanent disk error was collected from the operating system logs on the specified server, or all paths that are associated with the specified disk are disconnected.

Action

Check the operating system logs on the specified server for the following types of disk error: SC_DISK_ERR1, SC_DISK_ERR2, SC_DISK_ERR3, and SC_DISK_ERR6. Check the configuration of the disk. Ensure that at least one path that is associated with the specified disk is connected.

ALR1353W The status of disk *disk name* on server *server name* has degraded to Warning from *former status*.

Explanation

At least one temporary disk error was collected from the operating system logs on the specified server, or some paths that are associated with the specified disk are disconnected.

Action

Check the operating system logs on the specified server for the SC_DISK_ERR4 disk error. Check the configuration of the disk. Ensure that all paths that are associated with the specified disk are connected.

ALR1354I The status of disk *disk name* on server *server name* has improved to Normal from *former status*.

Explanation

No errors were collected from the operating system logs on the specified server. All paths that are associated with the specified disk are connected.

Action

No action is required.

ALR4625I New *entity type*, *entity name*, added to system *type system name* .

Explanation

New entity added.

Action

No action is required.

BPCCA - Data collector installation messages

- [BPCCA0001I](#) The data collector started, connected to the storage management service, and is ready to process requests from the storage management service.
- [BPCCA0002E](#) The data collector failed to connect to the storage management service at server_url.
- [BPCCA0003E](#) The data collector started but detected a problem with the directory directory_name and must stop.
- [BPCCA0004E](#) The data collector cannot run because it is not configured correctly.
- [BPCCA0005E](#) The data collector failed to connect to the storage management service at server_url because the host name could not be resolved.
- [BPCCA0006E](#) The data collector failed to connect to the storage management service at server_url because of an unknown error.
- [BPCCA0007E](#) The data collector failed to connect to a service from the storage management system.
- [BPCCA0008E](#) The data collector failed to connect to the storage management service because of invalid credentials.
- [BPCCA0009I](#) The data collector connected to the storage management service. The data collector had failed to connect since date_and_time.
- [BPCCA0010E](#) The data collector in the directory_name directory of the host host_name was running and an attempt was made to start a second instance of the same data collector. The second instance of the data collector stopped.
- [BPCCA0011I](#) The data collector stopped because a user requested it to shut down.
- [BPCCA0012I](#) The data collector stopped to enable the installation of an upgraded version of the data collector.
- [BPCCA0013E](#) The storage management service did not allow the data collector to connect because another data collector was already connected to the service.
- [BPCCA0100I](#) The updateCollector utility started.
- [BPCCA0101E](#) The collector directory was not specified in the collectorDirectory.properties file.
- [BPCCA0102E](#) The collector directory directory_name that was specified in the collectorDirectory.properties file is invalid. The collector directory is the directory to which the updateCollector utility must copy the upgrade image files.
- [BPCCA0103E](#) The collector directory directory_name that was specified in the collectorDirectory.properties file cannot be used as the collector directory.
- [BPCCA0104E](#) The updateCollector utility started but there was a problem with the upgrade image directory current_directory.
- [BPCCA0105E](#) The collector directory directory_name that was specified in the collectorDirectory.properties file is a subdirectory of the upgrade image directory upgrade_image_directory_name.
- [BPCCA0106E](#) Cannot upgrade the data collector in the collector_directory_directory because the directory contains the following locked files: locked_files_list

- [BPCCA0107I](#) The content of the `directory_collector_directory` will be deleted and replaced with subdirectories and files from the `upgrade_image_directory`. Some configuration files, the `log_directory`, and the contents of the `log_directory` will not be deleted.
- [BPCCA0108E](#) The data collector service cannot be uninstalled from the operating system. The upgrade process cannot be completed.
- [BPCCA0109E](#) The `updateCollector` utility could not upgrade the data collector. The data collector service is now in an inconsistent state.
- [BPCCA0110E](#) The contents of the `collector_directory` could not be deleted. The upgrade process cannot be completed. The data collector might be in an inconsistent state.
- [BPCCA0111E](#) The files and directories of the data collector from the `directory_upgrade_image_directory` could not be copied into the `directory_collector_directory`. The upgrade process cannot be continued. The data collector might be in an inconsistent state.
- [BPCCA0112I](#) The data collector in the `directory_collector_directory` was upgraded successfully to version `downloaded_version`.
- [BPCCA0113E](#) The data collector in the `directory_collector_directory` could not be upgraded.
- [BPCCA0114I](#) The data collector was upgraded to the new version and will start automatically.
- [BPCCA0115I](#) The attempt to upgrade the data collector failed. The existing data collector will start automatically.
- [BPCCA0116E](#) The attempt to upgrade the data collector failed. You must download and install the latest version of the data collector.
- [BPCCA0117E](#) The upgraded data collector could not be installed as a service on the operating system. You must install the new data collector service manually.
- [BPCCA0118I](#) The data collector was upgraded to the new version and started successfully.
- [BPCCA0119I](#) The data collector could not be upgraded, but was not modified. The existing data collector was restarted successfully.
- [BPCCA0120E](#) The upgraded data collector did not start.
- [BPCCA0121E](#) The existing data collector did not restart.
- [BPCCA0122E](#) The data collector cannot authenticate to the HTTPS proxy server `proxy_server_hostname`.

BPCCA0001I The data collector started, connected to the storage management service, and is ready to process requests from the storage management service.

Explanation

The data collector is running and ready to process requests from the storage management service.

Action

No action is required.

BPCCA0002E The data collector failed to connect to the storage management service at `server_url`.

Explanation

The data collector failed to connect to the storage management service. This condition might occur if you do not have a network connection to the service or if a firewall is preventing network access to the service. The condition might also occur if the service is unavailable.

Action

Verify that you have a network connection to the server. Verify that a firewall is not preventing network access to product services and agents. Verify that the storage management service is available.

BPCCA0003E The data collector started but detected a problem with the directory `directory_name` and must stop.

Explanation

The data collector started in the directory where the downloaded collector image was extracted. However, the structure of the directory is damaged. For example, some of the required subdirectories such as `bin`, `jre`, or `lib` might have been deleted.

Action

Download and extract the data collector image again.

BPCCA0004E The data collector cannot run because it is not configured correctly.

Explanation

The data collector is not configured correctly. The configuration file was not loaded, or at least one required configuration parameter was missing or had an invalid value.

Action

Check the trace_0.log file in the log directory of the data collector. Check the error message text that is associated with the InitializationException message. Check the other log files to try to determine the problem.

BPCCA0005E The data collector failed to connect to the storage management service at `server_url` because the host name could not be resolved.

Explanation

The host name could not be resolved by a name server.

Action

Verify that the host name for the storage management service can be resolved from the computer where the data collector is running. Verify that the name server is configured correctly.

BPCCA0006E The data collector failed to connect to the storage management service at `server_url` because of an unknown error.

Explanation

An error occurred while connecting to the storage management service, but the root cause cannot be determined.

Action

Verify that the local area network is available and that a firewall is not preventing network access. If you still cannot resolve the error, go to the IBM Support page (<https://www.ibm.com/mysupport/>) where you can chat with an expert, browse troubleshooting topics and forums, open support cases, and access IBM Documentation.

BPCCA0007E The data collector failed to connect to a service from the storage management system.

Explanation

A service from the storage management system failed to start or encountered a problem and stopped.

Action

Go to the IBM Support page (<https://www.ibm.com/mysupport/>) where you can chat with an expert, browse troubleshooting topics and forums, open support cases, and access IBM Documentation.

BPCCA0008E The data collector failed to connect to the storage management service because of invalid credentials.

Explanation

The credentials of the data collector were not authenticated by the storage management service.

Action

Verify that the data collector has the correct credentials. If you still cannot resolve the error, go to the IBM Support page (<https://www.ibm.com/mysupport/>) where you can chat with an expert, browse troubleshooting topics and forums, open support cases, and access IBM Documentation.

BPCCA0009I The data collector connected to the storage management service. The data collector had failed to connect since *date_and_time*.

Explanation

The data collector was able to connect to the storage management service, however it previously encountered connection errors.

Action

No action is required.

BPCCA0010E The data collector in the *directory_name* directory of the host *host_name* was running and an attempt was made to start a second instance of the same data collector. The second instance of the data collector stopped.

Explanation

Only one instance of the data collector can run in a directory at a time.

Action

No action is required because the second instance of the data collector stopped.

BPCCA0011I The data collector stopped because a user requested it to shut down.

Explanation

No further information is available.

Action

BPCCA0012I The data collector stopped to enable the installation of an upgraded version of the data collector.

Explanation

No further information is available.

Action

BPCCA0013E The storage management service did not allow the data collector to connect because another data collector was already connected to the service.

Explanation

Multiple data collectors cannot connect simultaneously to the same storage management service.

Action

If you want to use a new data collector instead of the data collector that is already connected, stop the data collector that is already connected. The new data collector will connect to the storage management service.

BPCCA0100I The `updateCollector` utility started.

Explanation

The data collector downloaded a new version of the data collector into an upgrade image directory. The `updateCollector` utility will update the data collector with the version from the upgrade image directory, and then restart the data collector.

Action

No action is required.

BPCCA0101E The collector directory was not specified in the `collectorDirectory.properties` file.

Explanation

The collector directory is the directory to which the `updateCollector` utility must copy the upgrade image files. The collector directory is specified in the `collectorDirectory.properties` file.

Action

Verify that the collector directory is specified in the `collectorDirectory.properties` file. Verify that the `updateCollector` utility was started by the data collector and not by any other means.

BPCCA0102E The collector directory `directory_name` that was specified in the `collectorDirectory.properties` file is invalid. The collector directory is the directory to which the `updateCollector` utility must copy the upgrade image files.

Explanation

The collector directory is the directory to which the `updateCollector` utility must copy the upgrade image files. The collector directory is specified in the `collectorDirectory.properties` file. The collector directory must be a directory that already exists and must not be the root directory.

Action

Verify that the directory that is specified in the `collectorDirectory.properties` file exists and is not the root directory.

BPCCA0103E The collector directory `directory_name` that was specified in the `collectorDirectory.properties` file cannot be used as the collector directory.

Explanation

The collector directory is the directory to which the `updateCollector` utility must copy the upgrade image files. The collector directory is specified in the `collectorDirectory.properties` file. The collector directory cannot be a directory that is reserved for use by the operating system.

Action

Verify that the directory that is specified in the `collectorDirectory.properties` file is not reserved for use by the operating system.

BPCCA0104E The `updateCollector` utility started but there was a problem with the upgrade image directory `current_directory`.

Explanation

The updateCollector utility must be started from the upgrade image directory, that is, the directory to which the upgraded collector was downloaded.

Action

Verify that the updateCollector utility started from the upgrade image directory.

BPCCA0105E The collector directory *directory_name* that was specified in the collectorDirectory.properties file is a subdirectory of the upgrade image directory *upgrade_image_directory_name*.

Explanation

The collector directory is the directory to which the updateCollector utility must copy the upgrade image files. The collector directory is specified in the collectorDirectory.properties file. The collector directory cannot be a subdirectory of the upgrade image directory.

Action

Verify that the directory that is specified in the collectorDirectory.properties file is not a subdirectory of the upgrade image directory.

BPCCA0106E Cannot upgrade the data collector in the *collector_directory* directory because the directory contains the following locked files: *locked_files_list*

Explanation

The directory for the data collector that you want to upgrade contains some locked files. The files might be in use by other processes. The files cannot be deleted or overwritten. The data collector cannot be upgraded. The data collector, which has not been upgraded, will restart automatically.

Action

After the data collector restarts, try to upgrade the collector again from a browser.

BPCCA0107I The content of the directory *collector_directory* will be deleted and replaced with subdirectories and files from the *upgrade_image_directory* directory. Some configuration files, the log directory, and the contents of the log directory will not be deleted.

Explanation

No further information is available.

Action

No action is required.

BPCCA0108E The data collector service cannot be uninstalled from the operating system. The upgrade process cannot be completed.

Explanation

The updateCollector utility cannot uninstall the data collector service from the operating system.

Action

Run the `uninstallDataCollectorService.sh` script from the collector directory to uninstall the data collector service. Delete the contents of the collector directory, then download and install the latest version of the data collector.

BPCCA0109E The `updateCollector` utility could not upgrade the data collector. The data collector service is now in an inconsistent state.

Explanation

The `updateCollector` utility could not uninstall the data collector service from the operating system.

Action

If the collector directory has not been deleted, run the `uninstallDataCollectorService.bat` script. Alternatively, use a command window to run the `uninstallDataCollectorService.bat` script from the upgrade image directory. Use the absolute path to the collector directory with the `uninstallDataCollectorService.bat` command from the upgrade image directory.

If you still cannot uninstall the data collector service, get the name of the service from the Services tab on the Windows Task Manager. Use the Windows Server "sc" command with the name of the data collector service to delete the service. For example, run the following command:

```
sc delete datacollector1
```

Delete the contents of the collector directory, then download and install the latest version of the data collector.

BPCCA0110E The contents of the `collector_directory` directory could not be deleted. The upgrade process cannot be completed. The data collector might be in an inconsistent state.

Explanation

No further information is available.

Action

Delete the data collector files from the collector directory, then download and install the latest version of the data collector.

BPCCA0111E The files and directories of the data collector from the directory `upgrade_image_directory` could not be copied into the directory `collector_directory`. The upgrade process cannot be continued. The data collector might be in an inconsistent state.

Explanation

No further information is available.

Action

Delete the data collector files from the collector directory, then download and install the latest version of the data collector.

BPCCA0112I The data collector in the directory `collector_directory` was upgraded successfully to version `downloaded_version`.

Explanation

No further information is available.

Action

No action is required.

BPCCA0113E The data collector in the directory *collector_directory* could not be upgraded.

Explanation

The log files in the log directory of the data collector might contain information about why the upgrade failed.

Action

Check the message.log and the trace_0.log files in the log directory of the data collector. The existing data collector will restart automatically. If the existing data collector does not restart, you must download and install the latest version of the data collector.

BPCCA0114I The data collector was upgraded to the new version and will start automatically.

Explanation

No further information is available.

Action

No action is required.

BPCCA0115I The attempt to upgrade the data collector failed. The existing data collector will start automatically.

Explanation

No further information is available.

Action

After the existing data collector is restarted, try again to download and install the latest version of the data collector.

BPCCA0116E The attempt to upgrade the data collector failed. You must download and install the latest version of the data collector.

Explanation

The attempt to upgrade the data collector failed. It is likely that the data collector is now in an inconsistent state.

Action

Delete the existing collector directory and its contents, then download and install the latest version of the data collector.

BPCCA0117E The upgraded data collector could not be installed as a service on the operating system. You must install the new data collector service manually.

Explanation

The data collector was upgraded to the new version but the upgraded data collector could not be installed as a service in the operating system. You must install the new data collector service manually.

Action

Log into the computer that the data collector and its file are on. On a UNIX system, run `installDataCollectorService.sh` as a root user. On a Windows system, run `installDataCollectorService.bat` as a Windows administrator. If other security restrictions are enabled, you might need to use the "Run as administrator" option to run `installDataCollectorService.bat`.

BPCCA0118I The data collector was upgraded to the new version and started successfully.

Explanation

The data collector was upgraded to the new version that was downloaded. The upgraded data collector started successfully.

Action

No action is required.

BPCCA0119I The data collector could not be upgraded, but was not modified. The existing data collector was restarted successfully.

Explanation

No further information is available.

Action

Try again to download and install the latest version of the data collector.

BPCCA0120E The upgraded data collector did not start.

Explanation

The data collector was upgraded, but did not start. You must start the data collector from the operating system.

Action

Log into the computer that the data collector and its file are on. On a UNIX system, run the `dataCollector.sh` script to start the data collector. On a Windows system, start the service from the Services tab on the Windows Task Manager.

BPCCA0121E The existing data collector did not restart.

Explanation

The data collector was not upgraded. The existing data collector was not modified. The `updateCollector` utility tried to restart the existing data collector, but failed.

Action

To restart the data collector, log into the computer that the data collector and its file are on. On a UNIX system, run the `dataCollector.sh` script to start the data collector. On a Windows system, start the service from the Services tab on the Windows Task Manager.

Try again to download and install the latest version of the data collector.

BPCCA0122E The data collector cannot authenticate to the HTTPS proxy server `proxy_server_hostname`.

Explanation

The data collector cannot connect to the storage management service.

Action

The authentication credentials for the proxy server might have changed. Verify that the data collector has the correct credentials for HTTPS proxy server.

Verify that the data collector is configured correctly to access the HTTPS proxy server.

BPCIN - Spectrum Control installation messages

- [BPCIN0001I The system is installing IBM Spectrum Control.](#)
- [BPCIN0002I The system completed the installation of IBM Spectrum Control.](#)
- [BPCIN0003E The installation program could not find the file file_name in the installation image. For more information, go to the IBM Knowledge Center and search on the message code.](#)
- [BPCIN0004E An error occurred during the installation of the component_name component. Review the log files in the log_file_directory_name directory for additional information.](#)
- [BPCIN0005E Invalid characters "characters" were found in the installation path "path".](#)
- [BPCIN0006E The location "location" that was specified for the installation is not empty. It might contain hidden items.](#)
- [BPCIN0007E Directory directory is not writable.](#)
- [BPCIN0008E The installation location specified is blank. Enter an installation location.](#)
- [BPCIN0009E The host name or IP address is not valid for name_with_spaces.](#)
- [BPCIN0010E The port number field is blank. Enter a valid port number.](#)
- [BPCIN0011E Enter a fully qualified host name or IP address.](#)
- [BPCIN0012E Enter a valid port number.](#)
- [BPCIN0013E Enter a port number from 1 to 65535.](#)
- [BPCIN0014E The user name cannot be blank or contain spaces. Enter a valid user name.](#)
- [BPCIN0015E The location for the license key file is blank. Enter a file location.](#)
- [BPCIN0016E The file that was specified does not exist. Enter a valid license key file.](#)
- [BPCIN0017E The file that was specified is not a valid license key file.](#)
- [BPCIN0018E The installation path cannot end with the characters "endCharacter".](#)
- [BPCIN0019I The system is installing the component component.](#)
- [BPCIN0020I The system completed the installation of the component component.](#)
- [BPCIN0021E An unexpected error occurred during the prevalidation of the component.](#)
- [BPCIN0022E The port range port_start - port_end is not available because the port or ports port are already used by other applications. The next available port range is available_start_port - available_end_port.](#)
- [BPCIN0023E The password is incorrect. It cannot be blank or contain spaces. Enter a valid password.](#)
- [BPCIN0024E The user name userID cannot contain the following special characters: characters.](#)
- [BPCIN0025E The password password cannot contain spaces or any of the following special characters: characters.](#)
- [BPCIN0026E The installation location selected does not have enough space.](#)
- [BPCIN0027E The user name userID is not in the administrative group adminGroup.](#)
- [BPCIN0028E The validation for user name userID has failed. Check to see if this user name exists.](#)
- [BPCIN0029E The password that was entered does not match the password for user_name on the system. Please try again.](#)
- [BPCIN0031E The configuration of component failed. Review the installation log file.](#)
- [BPCIN0032E The host name field is empty. Enter the fully qualified host name or IP address.](#)
- [BPCIN0033E The specified host name or IP address cannot be identified.](#)
- [BPCIN0034E The installation location is not an absolute path. Enter an absolute path.](#)
- [BPCIN0035W The installation program cannot validate the host name because the fully qualified domain name \(FQDN\) cannot be retrieved for the host.](#)
- [BPCIN0043E The last port number must be a port number from 1 to 65535.](#)
- [BPCIN0044E The installation location that was specified uses a Windows reserved name. Enter a different installation location.](#)
- [BPCIN0045E The installation location cannot contain special shell characters special_characters in the installation path path.](#)
- [BPCIN0046E The installation location contains special characters special_characters that are not supported by the operating system in the installation path path.](#)
- [BPCIN0047E The IBM Db2 database manager must be active to continue the installation process.](#)
- [BPCIN0048E IBM Spectrum Control cannot be installed because Db2 is not installed on the system or the Db2 configuration is not valid.](#)
- [BPCIN0049E The Db2 db2Version is not supported. The minimum supported Db2 versions are minDB2Version.](#)
- [BPCIN0050E An error occurred when the installation program tried to verify that the Db2 database manager is running. The error message is: error_message.](#)
- [BPCIN0051E An error occurred when the installation program tried to find an available database named dbName.](#)
- [BPCIN0052I The name of the database to be created is dbName.](#)
- [BPCIN0053E An error occurred when creating the database dbName. Review the log files for more information.](#)
- [BPCIN0055E Db2 is not installed or the Db2 profile was not sourced before installing IBM Spectrum Control.](#)
- [BPCIN0056E The user name userID does not have write permission on the default database path configuration parameter DFTDBPATH: dftdbpath.](#)
- [BPCIN0057E The user name userID is not in an operating system group that has Db2 SYSADM authority. Before you run your DB2 installation, validate that the Windows Server Hostname is 15 characters or less in length.](#)
- [BPCIN0059W There are long file names in the IBM Spectrum Control installation images.](#)
- [BPCIN0060E An error occurred during the uninstallation of component. Review the log files for more information.](#)
- [BPCIN0061E An invalid host name or IP address was specified for the Data server.](#)
- [BPCIN0062E The IPv6 internet protocol is not enabled on the specified host computer.](#)
- [BPCIN0063E The Data server is not running. For more information, go to the IBM Knowledge Center and search on the message code.](#)
- [BPCIN0064E The Data Server is not running at the specified host address or port: host host_address, port port. For more information, go to the IBM Knowledge Center and search for the message code.](#)
- [BPCIN0066E Errors occurred during the installation of the IBM Spectrum Control GUI. Review the log files for more information.](#)
- [BPCIN0068E Errors occurred during the uninstallation of the IBM Spectrum Control GUI. Review the log files for more information.](#)
- [BPCIN0069E Errors occurred during the configuration of Tivoli Common Reporting for IBM Spectrum Control. Review the log files for more information.](#)
- [BPCIN0070E Errors occurred during the configuration of the IBM Spectrum Control data model in Tivoli Common Reporting. Review the log files for more information.](#)
- [BPCIN0071E An error occurred because the installation program could not find the Db2 DFTDBPATH variable. For more information, go to the IBM Knowledge Center and search for the message code.](#)
- [BPCIN0072E The validation for user name userID has failed. Check to see if this user name exists or if Db2 is running.](#)

- [BPCIN0074E](#) The port range validation failed because the port value is not numeric.
- [BPCIN0075E](#) An invalid GUID 0xFFFFFFFFFFFFFFFF was found. Update or uninstall the GUID.
- [BPCIN0076E](#) IBM Spectrum Control could not read the GUID. See the GUID installation log for an explanation of the error.
- [BPCIN0077E](#) The name specified for the database is not valid because it contains a space or is blank.
- [BPCIN0078E](#) The database name database_name is not valid. The name can only contain the following characters: a-z, A-Z or 0-9.
- [BPCIN0079E](#) The first character or characters in the database name are not valid. The name must not begin with a number or the letters: SYS, DBM, or IBM.
- [BPCIN0080E](#) The database name specified is not valid because an existing database has the same name: database_name.
- [BPCIN0081E](#) The database name database_name is too long. The name can be 1 - 8 characters in length with no spaces.
- [BPCIN0082E](#) The specified host name or IP address is not a remote host. Enter a remote host name or IP address.
- [BPCIN0083E](#) The server could not connect to the remote database. Verify that the port is correct and that the database is running on the remote server. Also verify that the internet protocol connection between the server and remote database is compatible.
- [BPCIN0084E](#) The database name is not valid because there is no database on the specified server with this name.
- [BPCIN0085E](#) The installation program was unable to connect to the remote database. Check the remote server to verify that you can connect to Db2. For more information about Db2 connection issues, go to the IBM Knowledge Center at <http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome>.
- [BPCIN0086E](#) The connection to the remote database failed because the user name or password is not valid.
- [BPCIN0087E](#) The remote host name or IP address is not valid.
- [BPCIN0088E](#) The database repository was not found on the remote server.
- [BPCIN0089E](#) The version "retrivedVersion" of the remote database repository is not at the correct level. Upgrade the remote database repository to version "requiredVersion", then upgrade the remaining components.
- [BPCIN0090E](#) The name specified for the database is not valid because it contains a space.
- [BPCIN0091E](#) The file path file_path specified is not an absolute path or is not a directory or the partition does not exists.
- [BPCIN0092E](#) The first directory directory_name specified is not valid.
- [BPCIN0093E](#) The path string path_string is too long. The path string cannot be longer than 242 bytes.
- [BPCIN0094E](#) The user name user_name does not have write permission on the specified database path: database_path.
- [BPCIN0095E](#) The user name user_name does not have write permission on the specified log location: log_location.
- [BPCIN0096E](#) The database path cannot be blank. Enter a valid database path.
- [BPCIN0097E](#) The log location cannot be blank. Enter a valid log location.
- [BPCIN0098E](#) Enter 10 or fewer paths.
- [BPCIN0099E](#) You have entered pathNumber directories. A maximum of 10 directories can be specified.
- [BPCIN0100W](#) The database log files and database are in the same location. Click Yes to change the database log files or database location. Click No to ignore this message.
- [BPCIN0101E](#) The location "location" that was specified contains Db2 logs or log files.
- [BPCIN0102E](#) The Db2 profile for the "db2_instance" instance was not run before installing IBM Spectrum Control.
- [BPCIN0103E](#) The user name userID is not in the system group adminGroup.
- [BPCIN0104I](#) Resuming a failed install and installing the remaining components.
- [BPCIN0105E](#) A reboot must be done before continuing with the installation.
- [BPCIN0106E](#) There are invalid header files in the installation images. You may might be using the AIX tar program instead of the GNU tar program to extract files from the installation images. For more information, see IBM Knowledge Center and search by the error message code.
- [BPCIN0107E](#) The database path exists in the database path list. You must specify a unique database path.
- [BPCIN0108E](#) An error occurred during the upgrade of the component. Review the log files in the following directory for an explanation of the error: location.
- [BPCIN0109E](#) An unexpected error occurred. IBM Spectrum Control cannot resolve this error. For more information, review the log files and go to the IBM Knowledge Center.
- [BPCIN0112I](#) Upgrading the components:upgrade_components.
- [BPCIN0116I](#) The system is upgrading the component component.
- [BPCIN0117I](#) The system completed the upgrade of the component component.
- [BPCIN0120E](#) The validation of user name and password could not be completed
- [BPCIN0121E](#) The installation type entered is not supported.
- [BPCIN0122E](#) The TPCCommon.dll or libTPCCommon.so library could not be found or loaded.
- [BPCIN0123E](#) The user has no administrative rights to install or uninstall IBM Spectrum Control.
- [BPCIN0124E](#) The operating system cannot find the Db2 Service name.
- [BPCIN0125E](#) The installation program cannot find the Db2 instance name. Check the Db2 instance configuration.
- [BPCIN0126E](#) Unable to find Db2 installation path. Check your Db2 configuration.
- [BPCIN0127E](#) The service name configuration parameter cannot be queried from the Db2 Database Manager.
- [BPCIN0129E](#) The version that is installed cannot be upgraded because it does not meet the minimum build version for the database repository. The minimum build version supported is build_version.
- [BPCIN0130E](#) The database administrator password in IBM Spectrum Control does not match the database administrator password for Db2. Run the changepasswords tool to update the database administrator password.
- [BPCIN0131E](#) An invalid host name or IP address was specified for the Device Server.
- [BPCIN0132E](#) The Device Server is not running at the specified host address or port: host_address.port port.
- [BPCIN0134E](#) The installation program could not stop the . Please manually stop the .
- [BPCIN0135E](#) Tivoli Storage Productivity Center version is installed. Before you can upgrade to version 5, you must upgrade Tivoli Storage Productivity Center to version 4.
- [BPCIN0136E](#) The version "retrivedVersion" of the remote database repository is not at the correct build level. Upgrade the remote database repository to the correct build level, then upgrade the remaining version "requiredVersion" components.
- [BPCIN0137E](#) The version "retrivedVersion" of the remote database repository is not at the correct build level. Ensure that the installation image for the remaining version "requiredVersion" components is at the same build level as the remote database repository, then upgrade the remaining components.
- [BPCIN0138E](#) The common user name password does not match the password for the existing user name user_name on the system. Run the changepasswords tool to change the common user name password for IBM Spectrum Control.
- [BPCIN0139I](#) The User Migration Tool completed successfully.
- [BPCIN0140E](#) The User Migration Tool could not be started. Start the User Migration Tool from the graphical user interface after the upgrade is finished.
- [BPCIN0143E](#) Errors occurred during the upgrade of the GUI. Review the log files for more information.
- [BPCIN0146E](#) The installation language specified is not a supported language.
- [BPCIN0148E](#) IBM Spectrum Control cannot be installed, because the physical memory size on this computer is too small. The minimum memory size is minMemoryProduction. Increase the physical memory on this computer and run the installation program again.
- [BPCIN0149W](#) The physical memory size on this computer is below the minimum requirements that are specified for a production system. The minimum memory size for a production system is minMemoryProduction. You must increase the physical memory on this computer and run the installation program again. If you install IBM Spectrum Control with lower memory, you can only use it in an evaluation environment.
- [BPCIN0150E](#) The installation program cannot determine the amount of physical memory on the system.
- [BPCIN0151E](#) You tried to install IBM Spectrum Control on an unsupported operating system. For more information on the supported operating systems, go to following link.

- [BPCIN0152E IBM Spectrum Control cannot determine if the operating system is supported.](#)
- [BPCIN0153E The installation program could not rename the jre folder. To continue the installation, stop all Java processes that access the jre directory: jre_folder.](#)
- [BPCIN0159E The web server data source creation failed.](#)
- [BPCIN0163E The service name configuration parameter cannot be queried from the Db2 Database Manager. Reboot the machine after Db2 installation. SQL Error Message is sqlerrmsg](#)
- [BPCIN0164E The service name configuration parameter cannot be queried from the Db2 Database Manager. SQL Error Message is sqlerrmsg](#)
- [BPCIN0165E The password for user name has expired. You must change the password, or select another user name, to install IBM Spectrum Control.](#)
- [BPCIN0166W The password for user name cannot be checked for expiration. Please ensure that it is not expired.](#)
- [BPCIN0167E You cannot upgrade the license to the same or lower level.](#)
- [BPCIN0168E An error occurred while checking the user name and password. Review the log files for more information.](#)
- [BPCIN0169E The location "location" that should be specified for the license is not correct.](#)
- [BPCIN0170E IBM Spectrum Control cannot upgrade the license because the upgrade must be done on the server system.](#)
- [BPCIN0173E An error occurred during the deployment of the file_name file.](#)
- [BPCIN0176E An error occurred during Db2 catalog creation, and the catalog was not created.](#)
- [BPCIN0177E You cannot upgrade IBM Spectrum Control with a license key file that is at a lower level than the installed license key file.](#)
- [BPCIN0178E The provided license key file_name is invalid.](#)
- [BPCIN0179E An error occurred during Db2 catalog deletion, and the catalog was not removed.](#)
- [BPCIN0181E The web server data source testing has failed.](#)
- [BPCIN0182E TPC-GUI war data source testing failed.](#)
- [BPCIN0184E The remote database repository already contains data. Install a new IBM Spectrum Control remote database repository first, then install the remaining IBM Spectrum Control components.](#)
- [BPCIN0185E The remote database repository for version oldVersion is incompatible with IBM Spectrum Control Version requiredVersion. Upgrade remote database repository to Version requiredVersion first, and then install the remaining IBM Spectrum Control components.](#)
- [BPCIN0190E You cannot use non-standard characters, such as a space between characters or an underscore, in a host name. You must enter a host name with standard characters and try again.](#)
- [BPCIN0191E You cannot install Db2 in a directory that starts with the letter a.](#)
- [BPCIN0193E The version "retrivedVersion" of the remote database repository is not at the correct level. Install a new Version "requiredVersion" remote database repository at the correct level, then install the remaining IBM Spectrum Control components.](#)
- [BPCIN0195E The version "retrivedVersion" of the database repository dbHost is not at the correct build level. Install a new Version "requiredVersion" database repository at the correct build level, then install the remaining IBM Spectrum Control components.](#)
- [BPCIN0198E The path to installation image "installDirectory" contains the invalid character "unallowedCharacter". You must change the directory name so that it contains valid characters.](#)
- [BPCIN0199E The installation program cannot validate the host name because the fully qualified domain name \(FQDN\) cannot be retrieved for the host.](#)
- [BPCIN0200E The installation program does not allow host name specified as IP address. Please ensure that a fully qualified domain name \(FQDN\) is provided.](#)
- [BPCIN0202E Runtime errors have occurred during the IBM Spectrum Control preinstallation process. The installation program cannot recover from this error.](#)
- [BPCIN0203E The installation program does not support backslash character in the user name. In case it is a Windows domain account, please specify just the user name without using the "Domain_name\\" prefix.](#)
- [BPCIN0205E The IBM Spectrum Control installation program could not find the directory "installDirectory" on the installation image. For more information, go to the IBM Knowledge Center and search on the message code.](#)
- [BPCIN0206E The IBM Spectrum Control installation program was unable to retrieve a fully qualified domain name \(FQDN\) for the host. You must configure the host system with an FQDN.](#)
- [BPCIN0207E You cannot upgrade when the stand-alone GUI is running. Stop the stand-alone GUI and click OK to continue the upgrade or click Quit to exit the installation program.](#)
- [BPCIN0208E The fully qualified domain name \(FQDN\) retrieved for the host contains non-standard characters, such as a space between characters or an underscore. You must configure the host system with an FQDN that contains standard characters.](#)
- [BPCIN0209E The installation program does not allow a user that is present in both the Windows Domain and the local Operating System repositories. Install with a user name that is only present in one of these repositories.](#)
- [BPCIN0210E The Db2 "sourcedb2profile" profile has not been loaded in the .profile file for user commonUser. This profile must be loaded before you install IBM Spectrum Control.](#)
- [BPCIN0211E The technology level or the service pack level of this operating system is not supported and must be upgraded to a supported level. The detected operating system version is os_version](#)
- [BPCIN0214E Cannot complete a fresh install without a valid license.](#)
- [BPCIN0215E The path extractorDirectory where IBM Spectrum Control is extracted has extractorDirLength characters and the maximum number of characters allowed is 260. You must shorten this path.](#)
- [BPCIN0217E The Windows registry check indicates that .NET 3.5 or higher is not available on this Windows server. You must install .NET 3.5 or higher to continue. To install .NET 3.5 type the following commands in a 64 bit windows powershell: Import-Module ServerManager Add-WindowsFeature as-net-framework](#)
- [BPCIN0219E The domain configuration is invalid. To resolve the issue, complete the following steps: Disable the Windows Firewall service. Start or restart the Computer Browser service on this domain member computer and on the domain controller computer. If the service has a Stopped or Disabled status on the domain controller computer, you must restart Computer Browser service on the domain member computer after you start the service on the domain controller computer. In a Windows command window, run the net view command and verify that there are no errors. Reinstall IBM Spectrum Control.](#)
- [BPCIN0220E The current login user loginuser is not an administrator or a member of the domain administrator group. The installation program cannot start.](#)
- [BPCIN0221E The user name userID is not a part of the local administrator group.](#)
- [BPCIN0222E The current login user loginuser is not an administrator or a member of the local administrator group. The installation program cannot continue.](#)
- [BPCIN0223E The current login user loginuser is not a member of the local Db2 administrator group. The installation program cannot start.](#)
- [BPCIN0224E The user userID is a domain account and cannot be used to log in to Db2. You must enter a separate user name for Db2.](#)
- [BPCIN0225E The current logged in user loginuser does not have Db2 SYSADM authority. To provide the user with Db2 SYSADM authority, log in by using a user name with SYSADM authority and run the following commands: db2cmd db2set -g DB2_GRP_LOOKUP=local,TOKENLOCAL db2 force application all db2stop db2start](#)
- [BPCIN0226E Login as a windows domain user in order to install IBM Spectrum Control using a windows domain account.](#)
- [BPCIN0229E The Monitoring Agent for Windows OS - Watchdog and Monitoring Agent for Windows OS - Primary services must be stopped before you can continue. After you upgrade, you must restart these services.](#)
- [BPCIN0230E The Monitoring Agent for Windows OS - Watchdog and Monitoring Agent for Windows OS - Primary services must be stopped before you can continue. After you uninstall IBM Spectrum Control, you must restart these services if the required reboot is postponed.](#)
- [BPCIN0231E Tivoli Common Reporting is not installed on your system. You must install it to continue with the IBM Spectrum Control installation.](#)

- [BPCIN0233E Jazz for Service Management is not installed in the specified location. Reenter the correct installation location for Jazz for Service Management.](#)
- [BPCIN0234E The Jazz for Service Management user credentials that you entered were incorrect. Check the user credentials and try again.](#)
- [BPCIN0235W The Jazz for Service Management and Tivoli Common Reporting servers are not running. To start the servers, run the following command: file It takes a few minutes for these servers to start up and initialize. Wait for a few minutes before resuming the installation.](#)
- [BPCIN0236W The Tivoli Common Reporting server is not running. If you have already started the server, wait for a few more minutes. It takes a while for the server to start up and initialize. Otherwise, you can start the server by running this command: file](#)
- [BPCIN0238E The Tivoli Common Reporting server at tcrlocation cannot be reached. If the Jazz for Service Management server has not completed startup, wait for a few more minutes and then click OK. It takes a while for the server to start and initialize. Otherwise, if the Jazz for Service Management server is started but Tivoli Common Reporting still cannot be reached, restart the Jazz for Service Management server using the following commands: stopServer - username username -password password startServer](#)
- [BPCIN0239E The Tivoli Common Reporting server configuration cannot be exported.](#)
- [BPCIN0240E The Tivoli Common Reporting server configuration file "configuration file" was not created.](#)
- [BPCIN0241E The Tivoli Common Reporting configuration cannot be upgraded.](#)
- [BPCIN0242E The Windows registry check indicates that .NET 3.5 is not available on this Windows 2012 server. You must install .NET 3.5 to continue. To install .NET 3.5 type the following commands in a 64 bit windows powershell: Import-Module ServerManager Add-WindowsFeature as-net-framework](#)
- [BPCIN0244E An error occurred while enumerating the local administrator group membership. On the current computer, on the Properties page of this group, remove the user names that are displayed with a security ID \(SID\). An example of a SID is S-1-5-21-337177553-1671989427-887411491-500 and is used instead of a user name.](#)
- [BPCIN0245E An error occurred while enumerating the local Db2 administrator group membership. On the current computer, on the Properties page of this group, remove the user names that are displayed with a security ID \(SID\). An example of a SID is S-1-5-21-337177553-1671989427-887411491-500 and is used instead of a user name.](#)
- [BPCIN0246E Tivoli Storage Productivity Center version is installed. Before you can upgrade to version 5, you must upgrade Tivoli Storage Productivity Center to version 4.](#)
- [BPCIN0247E An error occurred during the domain check prevalidation process. Verify that the domain controller computer is available and then restart this domain member machine.](#)
- [BPCIN0248E The IBM Spectrum Control installation program supports only fully qualified user names on Windows domain member machines. Specify the user name userID by using the "Domain_name\" or the "Machine_name\" prefix. The detected domain name is "domainName".](#)
- [BPCIN0250W Before you can upgrade the database repository, you must first stop the Data server, Device server, Alert server, Export server, and Web GUI server on the remote server. For more information about stopping IBM Spectrum Control servers use the following link.](#)
- [BPCIN0252E Db2 has been installed by using a domain user account. The IBM Spectrum Control installation software does not allow the Db2 user name userID because this user name exists in both the Windows domain and the local operating system repositories. Specify a Db2 user name that exists only in the Windows Domain registry.](#)
- [BPCIN0253W All IBM Spectrum Control components are installed. GUIModeMessage](#)
- [BPCIN0254W IBM Spectrum Control servers are installed on remoteHost.](#)
- [BPCIN0255E The database dbName must be upgraded to the current version. For more information about upgrading Db2, go to the IBM Knowledge Center and search for "Upgrading Db2"](#)
- [BPCIN0256W If you continue upgrading to IBM Spectrum Control, the reports you have from Tivoli Storage Productivity Center Version oldVersion will be deleted. When you uninstall Tivoli Integrated Portal, the Authentication Services Server is also uninstalled. If you have storage subsystems that are configured to use LDAP authentication through the Authentication Services Server, before you upgrade the product and uninstall Tivoli Integrated Portal, reconfigure the storage subsystems so that these subsystems do not use LDAP authentication. Click OK to continue or Cancel to select another option.](#)
- [BPCIN0257E The Db2 database installation on dbPath path was made when the creation of 8.3 filenames was disabled on this server. The IBM Spectrum Control installation will fail when Db2 is installed to a path with spaces and 8.3 filenames were not enabled when it was installed.](#)
- [BPCIN0258E The LDAP configuration export did not succeed. For more information, review the log files, and in the LDAP export command output, search for "exportLDAPRepositories".](#)
- [BPCIN0260E The connection to the local database failed because the user name or password is not valid.](#)
- [BPCIN0261E The installation program was unable to connect to the local database. Check the server to verify that you can connect to Db2. For more information about Db2 connection issues, go to the IBM Knowledge Center at <http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome>.](#)
- [BPCIN0262E The IBM Spectrum Control upgrade process has stopped because Jazz for Service Management and Tivoli Integrated Portal are using the same ports. To continue, you must install Jazz for Service Management and Tivoli Common Reporting again and use ports that are different from the ports that are used by Tivoli Integrated Portal.](#)
- [BPCIN0264E The version of the database repository component cannot be queried from the Db2 Database Manager. SQL Error Message is sqlerrmsg](#)
- [BPCIN0265E An error occurred during the prevalidation of the component because the file_name is missing.](#)
- [BPCIN0266E Db2 has been installed by using a domain user name. The user name dbuser is not a member of the domain Db2 administrator group domainDB2AdminGroup, so the IBM Spectrum Control installation program cannot start.](#)
- [BPCIN0267E The password password cannot start with the following special characters: characters.](#)
- [BPCIN0268E The installation program could not stop the . You must manually stop the by running the script.](#)
- [BPCIN0269E The Jazz\(tm\) for Service Management installation image Version "retrivedVersion" is not at the correct build level.](#)
- [BPCIN0272E "installDirectory" cannot be found where the Jazz\(tm\) for Service Management installation files are extracted. Extract the Jazz\(tm\) for Service Management installation files to a local directory before you start the Jazz\(tm\) for Service Management installation program.](#)
- [BPCIN0275E A 32-bit version of Db2 has been detected. IBM Spectrum Control can only be installed with a 64-bit version of Db2. Install a supported 64-bit version of Db2, and install IBM Spectrum Control again.](#)
- [BPCIN0276E A required system library could not be loaded. Review the log files for more information about this error.](#)
- [BPCIN0277W The Jazz\(tm\) for Service Management installation image that you selected is not the latest version. Click Yes to continue with the current version and not install or upgrade reports. Click No to install or upgrade to Jazz\(tm\) for Service Management Version "latest_JazzSM_available_version" and Tivoli Common Reporting Version "latest_TCR_available_version".](#)
- [BPCIN0279W IBM Spectrum Control does not support the version of Jazz\(tm\) for Service Management that is installed on this computer. You must upgrade to Jazz\(tm\) for Service Management minJazzSMVersion before you can continue.](#)
- [BPCIN0281E The current version of the Tivoli Common Reporting installation image is not at the correct level.](#)
- [BPCIN0282E The installation package for Tivoli Common Reporting Version 3.1.0.1 cannot be found on your computer. To continue installing IBM Spectrum Control, download the installation package for Tivoli Common Reporting Version 3.1.0.1 in the same directory where you downloaded the installation package for Tivoli Common Reporting Version 3.1.0.2.](#)
- [BPCIN0283E You cannot upgrade "latest_version" to the lower version "lower_version".](#)
- [BPCIN0284E The Jazz for Service Management installation directory does not have execution rights. You can add execution rights to the directory by running the following command: chmod -R u+x "JazzSM_build_folder"](#)
- [BPCIN0285E The getDB2Inst.sh "user_name" command is displaying a "null" result because the Db2 profile was not sourced for user "user_name" or the environment for "user_name" has been corrupted.](#)
- [BPCIN0286W The operation of the storage resource agent is limited on the operating system of the local server. Total Disk Space and Available Disk Space on this server cannot be determined.](#)
- [BPCIN0287E Directory directory is located on a memory based file system \(RAM disk\) and cannot be used for installing IBM Spectrum Control.](#)

- [BPCIN0288E](#) The upgrade process detected an error with the previously installed version of the product. The installation directory of the previously installed version of the product contains corrupted files. Remove or fix the corrupted files and run the upgrade process again. See the `lax*-out.txt` and `lax*-err.txt` log files in the system temporary directory for details of the error.
- [BPCIN0289E](#) You must install Db2 before you can install Cognos BI Reports. Ensure that Db2 is already installed on the system and the db2profile is sourced.
- [BPCIN0290E](#) The IBM Spectrum Control installation image is corrupted. Extract the IBM Spectrum Control installation image again into an empty directory.
- [BPCIN0291E](#) You need to define 'localhost' in the 'hosts' file that is used by Jazz for Service Management and restart the host system.
- [BPCIN0293E](#) The IBM Spectrum Control installation program cannot find the license folder. Ensure the license folder is extracted into "extractorDir".
- [BPCIN0295E](#) Tivoli Storage Productivity Center old_version is installed. Before you upgrade to IBM Spectrum Control version new_version, you must upgrade to Tivoli Storage Productivity Center version 5.2.7.
- [BPCIN0296E](#) The installation program could not rename the short_name directory because the folder is in use by other processes. Quit the installation program and stop all processes that access the directory full_path.
- [BPCIN0297E](#) The user_name user does not have the full control permission for the folder_path folder.
- [BPCIN0298E](#) The user_name fenced user does not have full permissions for the directory_path database directory.
- [BPCIN0299E](#) The installation location "install_location" has a naming conflict with a file or folder named folder_path. Rename the folder_path file or folder.
- [BPCIN0300E](#) The value db2_variable_existing_value is not allowed for the Db2 variable DB2_LIMIT_FENCED_GROUP. Use the `db2set -g DB2_LIMIT_FENCED_GROUP=OFF` command to change the value of the variable to OFF.
- [BPCIN0301E](#) The versionFile file is empty. Extract IBM Spectrum Control again and make sure the version.txt file contains a valid build string.
- [BPCIN0302E](#) Db2 Advanced Enterprise Server Edition is not supported. Install Db2 Enterprise Server Edition and then proceed with your installation.
- [BPCIN0303W](#) Your storage environment has switches with obsolete data sources. Configure up-to-date data sources after you upgrade.
- [BPCIN0304E](#) You cannot upgrade from a Basic Edition license.
- [BPCIN0306W](#) The Storage Resource agent registered successfully with the Data server but some problems occurred after the registration. Review the log files in the SRA_log_name directory for additional information.
- [BPCIN0307W](#) You are currently connected to devices using a CIM interface. This upgrade might require new certificates to be generated on your CIM managed devices to continue monitoring them. To resolve this issue, go to <http://www.ibm.com/support/docview.wss?uid=swg21976237>
- [BPCIN0307E](#) The security certificates for the Web server have expired. Renew the security certificates and run the installation program again. See the `lax*-out.txt` and `lax*-err.txt` log files in the system temporary directory for details of the error and how to resolve the issue.
- [BPCIN0308W](#) Brocade switches that use SNMP services were detected. SNMP data sources are no longer used to manage Brocade fabrics and switches. Instead, you must use the embedded SMI agent in Brocade Network Advisor.
- [BPCIN0309E](#) The port configuration export did not succeed. For more information, review the log files, and in the port export command output, search for "exportCurrentPorts".
- [BPCIN0310E](#) Spectrum Control installation requires that the Administrators group in Windows is assigned "Debug programs" privilege. Check this setting, log out of and back into Windows, and try again.
- [BPCIN0311E](#) Spectrum Control installation requires the wmic command to work. Check this command, ensure the service "Windows Management Instrumentation" is running and try again.
- [BPCIN0312E](#) Spectrum Control installation requires that wmic command to work. Check this command, ensure service "Windows Management Instrumentation" is working, disable any antivirus and try again.
- [BPCIN0313E](#) Spectrum Control installation requires that wmic command to work. Check this command and the PATH environment variable and try again.
- [BPCIN0314E](#) Spectrum Control installation requires that wmic command to work. Check this command and try again.
- [BPCIN0315E](#) Spectrum Control installation requires that chcp command to work. Check this command and try again.
- [BPCIN0316E](#) The IBM Spectrum Control upgrade process does not support custom materialized query tables (MQTs). Remove any MQTs from the database and run the upgrade process again.
- [BPCIN0317E](#) The IBM Spectrum Control upgrade process does not support missing vendor information from the .com.zerog.registry.xml file. Add the missing information and run the upgrade process again. See the `lax*-out.txt` and `lax*-err.txt` log files to resolve the issue.
- [BPCIN0318E](#) The IBM Spectrum Control upgrade process does not support duplicated users user in IBM WebSphere Application Server related to the Web server. Remove the duplicate user from a federated repository other than the localOS default repository; run the upgrade process again.
- [BPCIN0319E](#) The currently installed version of AIX XL C/C++ RUNTIME is not supported. You must upgrade to a supported level. The version that was detected is version.
- [BPCIN0320E](#) PAM (Pluggable Authentication Modules) isn't installed on your system. Installing PAM would resolve the issue.
- [BPCIN0321E](#) An error occurred when restoring authorities before creating the database dbName. Review the log files for more information.
- [BPCIN0323E](#) The Data Server port has different values in the configuration files. Verify that the port value is correct in the `installDir\data/config/server.config` file, and the `installDir/config/InstallVariable.properties` file. Also, verify that the PORT_NUMBER column where SERVER_TYPE is serverType is correct in the T_RES_Server database table.
- [BPCIN0324W](#) The current certificate used to secure the connection of the IBM Spectrum Control Data server and the Storage Resource agents cannot be replaced with the latest certificate that provides higher security. After you complete the upgrade, you can create new certificates by following the information in the following link.
- [BPCIN0325E](#) IBM Tivoli Storage Productivity Center old_version is installed. Before you can upgrade to IBM Spectrum Control new_version, you must upgrade to IBM Spectrum Control 5.3.0 or later. For more information about upgrading to IBM Spectrum Control 5.4.0 or later, go to the following link:
- [BPCIN0328E](#) The IBM Spectrum Control installation requires the LD_LIBRARY_PATH environment variable be set. Set this environment variable and try the installation again. If you ran `setup.bin` from the IBM Spectrum Control installation directory with `sudo` and the LD_LIBRARY_PATH variable was properly set, then try the `sudo -E LD_LIBRARY_PATH=$LD_LIBRARY_PATH ./setup.bin` command.
- [BPCIN0329E](#) The IBM Spectrum Control installation requires that the LIBPATH environment variable be set. Set this environment variable and try the installation again.
- [BPCIN0330E](#) An SQL exception was created when querying the Db2 Database Manager. The SQL error message is: `sqlerrmsg`.
- [BPCIN0331E](#) The `tpcregFile` file is not valid because it contains a failed or partial upgrade to a previous IBM Spectrum Control Version `newTPCVersion`. Continue with the incomplete upgrade before you try to upgrade to the new IBM Spectrum Control Version `varNewUpgradeVersion`.
- [BPCIN0335E](#) The IBM Spectrum Control installation requires you install the `libstdc++` package on the AIX operating system. Download and install the `libstdc++` package and start the installation again. You can download the package here: <https://www.ibm.com/developerworks/aix/library/aix-toolbox/alpha.html>.
- [BPCIN0336E](#) The version of the `libstdc++` package you installed is lower than the minimum required Version `minVersion`. Download and upgrade the `libstdc++` package and start the installation again. You can download the package here: <https://www.ibm.com/developerworks/aix/library/aix-toolbox/alpha.html>.
- [BPCIN0337E](#) The componentServer password failed to validate. Check that you entered the password correctly and try again.
- [BPCIN0338E](#) The componentServers passwords failed to validate. Check that you entered the passwords correctly and try again.
- [BPCIN0339W](#) Your current Db2 instance is using a trial license. The license will expire on `expireDate`. Once your trial expires Db2 will not start. If you want to upgrade your Db2 license now, use the information in the following link.
- [BPCIN0340W](#) Your current Db2 instance is using a trial license. The license will expire on `expireDate`. Once your trial expires Db2 will not start. If you want to upgrade your Db2 license now, goto link.
- [BPCIN0343W](#) There are new certificate requirements that are strictly enforced for macOS Catalina users which might affect their ability to access the IBM Spectrum Control GUI. During an upgrade of IBM Spectrum Control, certificates self-signed by IBM Spectrum Control will be made compliant automatically.

However, if one or more of your certificates are not self-signed by IBM Spectrum Control, see the following links for more information. Validate that your certificates are compliant.

- [BPCIN0344E Your current Db2 instance trial license has expired. If you want to upgrade your Db2 license now, goto the following link.](#)
- [BPCIN0345E The user user is part of the "Deny access to this computer from the network" security policy setting on your computer. Either contact your administrator to have the user removed from the security policy or enter a different user. For more information on required user privileges in installation scenarios, go to the following link.](#)
- [BPCIN0346E Your current Db2 License type: Community is not supported for IBM Spectrum Control. To upgrade your Db2 license, go to the following link:](#)

BPCIN0001I The system is installing IBM Spectrum Control.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCIN0002I The system completed the installation of IBM Spectrum Control.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCIN0003E The installation program could not find the file *file_name* in the installation image. For more information, go to the IBM Knowledge Center and search on the message code.

Explanation

The installation program could not find the file or files in the installation image. The file or files were not found because of one of the following reasons:

- An error occurred during the installation image download.
- A file was deleted from the installation image.
- An error occurred using a network drive.

Action

If the installation image was downloaded, verify that the installation image was downloaded correctly.

Correct the problem and click Resume to continue the installation.

If you cannot resolve the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0004E An error occurred during the installation of the *component_name* component. Review the log files in the *log_file_directory_name* directory for additional information.

Explanation

The specified IBM Spectrum Control component was not installed.

Action

Review the log files in one of the following directories:

- For Windows operating systems: TPC_installation_directory\logs\
- For AIX or Linux operating systems: TPC_installation_directory/logs/

TPC_installation_directory is where IBM Spectrum Control is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

Resolve the issue and run the installation program again. If you cannot resolve the problem, contact IBM Software Support.

Related reference

-  [Getting support](#)

BPCIN0005E Invalid characters "*characters*" were found in the installation path "*path*".

Explanation

The path that was specified for the installation contains characters that are not valid file-name characters.

The following characters are not valid in the installation path name:

- Parentheses ()
- Semicolon ;
- Question mark ?
- Exclamation point !
- Asterisk *
- Dollar sign \$
- Vertical bar |
- Double quotation mark "
- Less than <
- Greater than >
- Ampersand &

Action

Remove the special characters from the installation path name.

BPCIN0006E The location "*location*" that was specified for the installation is not empty. It might contain hidden items.

Explanation

The directory that was specified as the installation path contains files or folders.

Action

Select a directory that does not contain files or folders.

BPCIN0007E Directory *directory* is not writable.

Explanation

The directory that was selected for the installation is a read-only directory.

Action

Select a directory that is writable.

BPCIN0008E The installation location specified is blank. Enter an installation location.

Explanation

You must specify an installation location.

Action

Enter a valid path for the installation.

BPCIN0009E The host name or IP address is not valid for *name_with_spaces*.

Explanation

The host name or IP address that was specified is not valid because it is empty or contains one or more spaces.

Action

Enter a valid host name or IP address.

BPCIN0010E The port number field is blank. Enter a valid port number.

Explanation

The port number that was specified is not valid because the field is blank.

Action

Enter a valid port number.

BPCIN0011E Enter a fully qualified host name or IP address.

Explanation

The host name or IP address format that was specified is not valid because it is not a fully qualified host name or valid IP address.

An example of a fully qualified host name is: myuser.ibm.com.

Action

Enter a fully qualified host name or valid IP address.

BPCIN0012E Enter a valid port number.

Explanation

The port number that was specified is not valid.

Action

Enter a valid port number.

BPCIN0013E Enter a port number from 1 to 65535.

Explanation

The port number that was specified is not valid because it does not fall in the range from 1 to 65535.

Action

Enter a port number from 1 and 65535.

BPCIN0014E The user name cannot be blank or contain spaces. Enter a valid user name.

Explanation

The user name that was specified is not valid because it contains a space or is blank.

Action

Enter a valid user name that does not contain spaces or is blank.

BPCIN0015E The location for the license key file is blank. Enter a file location.

Explanation

The location that was specified for the license key file is not valid because it is blank.

Action

Enter a valid location for the license key file.

BPCIN0016E The file that was specified does not exist. Enter a valid license key file.

Explanation

A valid license key file must be specified.

Action

Enter a valid license key file.

BPCIN0017E The file that was specified is not a valid license key file.

Explanation

A valid license key file must be specified.

Action

Enter a valid license key file.

BPCIN0018E The installation path cannot end with the characters "*endCharacter*".

Explanation

The installation path that was specified is not valid because it ends with one or more of the following special characters:

- Comma ,
- Period .
- Double quotation mark "

Action

Remove the special characters from the installation path.

BPCIN0019I The system is installing the *component* component.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCIN0020I The system completed the installation of the *component* component.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCIN0021E An unexpected error occurred during the prevalidation of the *component*.

Explanation

The IBM Spectrum Control installation stopped because an error occurred during prevalidation. IBM Spectrum Control cannot resolve this error.

Action

To resolve this error, go to the IBM Spectrum Control support site, search for document "1302800", and contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0022E The port range *port_start - port_end* is not available because the port or ports *port* are already used by other applications. The next available port range is *available_start_port - available_end_port*.

Explanation

One or more ports in the port range are not available because the port or ports are used by another application.

Action

Specify an available port range.

BPCIN0023E The password is incorrect. It cannot be blank or contain spaces. Enter a valid password.

Explanation

The password is not valid because it is blank or contains one or more spaces.

Action

Enter a valid password.

BPCIN0024E The user name *userID* cannot contain the following special characters: *characters*.

Explanation

The user name that was specified cannot contain special characters.

Action

Enter a valid user name.

BPCIN0025E The password *password* cannot contain spaces or any of the following special characters: *characters*.

Explanation

The password cannot contain special characters.

Action

Enter a valid password.

BPCIN0026E The installation location selected does not have enough space.

Explanation

The location that was selected does not have enough space for the installation.

Action

Select a different location with more space.

BPCIN0027E The user name *userID* is not in the administrative group *adminGroup*.

Explanation

The user name to install IBM Spectrum Control must be in the administrative group.

Action

Add the user name to the administrative group and continue with the installation.

BPCIN0028E The validation for user name *userID* has failed. Check to see if this user name exists.

Explanation

The validation for the specified user name has failed because the user name might not exist on the system.

Action

Check to see if the user name exists on the system. Correct the problem and run the installation program again.

BPCIN0029E The password that was entered does not match the password for user *user_name* on the system. Please try again.

Explanation

The password must match the user name that exists on the system before you can specify the password and user name for installation.

Action

Verify that the password and user name matches on the system and then continue with the installation.

BPCIN0031E The configuration of *component* failed. Review the installation log file.

Explanation

An error occurred while configuring the specified component.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: TPC_installation_directory\logs\
- For UNIX or Linux operating system: TPC_installation_directory/logs/

TPC_installation_directory is where IBM Spectrum Control is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

Correct the problem and run the installation program again.

BPCIN0032E The host name field is empty. Enter the fully qualified host name or IP address.

Explanation

The host name must be a fully qualified host name or IP address.

Action

Enter the fully qualified host name or IP address.

BPCIN0033E The specified host name or IP address cannot be identified.

Explanation

The specified host name or IP address is not the correct host name or IP address.

Action

Enter the correct host name or IP address.

BPCIN0034E The installation location is not an absolute path.
Enter an absolute path.

Explanation

The installation location must be an absolute path and not a relative path.

An example of an absolute path for Windows is: C:\Program Files\IBM\TPC.

An example of an absolute path for UNIX or Linux is: /opt/IBM/TPC.

Action

Enter an absolute path.

BPCIN0035W The installation program cannot validate the host name because the fully qualified domain name (FQDN) cannot be retrieved for the host.

Explanation

The host name cannot be validated because the local fully qualified domain name cannot be retrieved for the host. Either the host name is incorrect or the DNS is providing an invalid configuration.

Action

Review and correct the Domain Name System (DNS) configuration on the host system or specify a valid host name. Continue with the installation.

BPCIN0043E The last port number must be a port number from 1 to 65535.

Explanation

The last port number must be a port number from 1 and 65535.

Action

Specify a port number from 1 to 65535.

BPCIN0044E The installation location that was specified uses a Windows reserved name. Enter a different installation location.

Explanation

The installation location cannot be a Windows reserved name. For information about Windows reserved names, see the Microsoft website.

Action

Enter an installation location that does not specify a Windows reserved name and run the installation program again.

BPCIN0045E The installation location cannot contain special shell characters *special_characters* in the installation path *path*.

Explanation

The installation location cannot contain special shell characters in the installation path.

Action

Enter an installation location that does not contain special shell characters and run the installation program again.

BPCIN0046E The installation location contains special characters *special_characters* that are not supported by the operating system in the installation path *path*.

Explanation

The installation location contains special characters that are not supported by the operating system.

Action

Enter an installation location that does not contain special characters and run the installation program again.

BPCIN0047E The IBM Db2 database manager must be active to continue the installation process.

Explanation

The Db2 database manager must be active to continue the installation process.

Action

Ensure that the Db2 instance is running before installing IBM Spectrum Control. Enter the following command from a command prompt window: db2start. Run the installation program again.

BPCIN0048E IBM Spectrum Control cannot be installed because Db2 is not installed on the system or the Db2 configuration is not valid.

Explanation

Db2 must be installed before you can install IBM Spectrum Control.

Action

Install Db2 on the system and then run the IBM Spectrum Control installation program again.

BPCIN0049E The Db2 *db2Version* is not supported. The minimum supported Db2 versions are *minDB2Version*.

Explanation

Before you can install IBM Spectrum Control, a supported version of Db2 must be installed on the system.

Action

For information about the versions of Db2 that are supported, go to the IBM Knowledge Center and search for "software requirements database repository". Install or upgrade the required Db2 version and run the installation program again.

BPCIN0050E An error occurred when the installation program tried to verify that the Db2 database manager is running. The error message is: *error_message*.

Explanation

An error occurred when IBM Spectrum Control tried to verify that the Db2 database manager is running. It is possible that the shell environment is not set up correctly for Db2.

Action

Go to the IBM Knowledge Center and search on the text string "Starting Installation". Follow the instructions on sourcing the Db2 profile then run the installation program again.

BPCIN0051E An error occurred when the installation program tried to find an available database named *dbName*.

Explanation

The installation program cannot recover from this error.

Action

Contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0052I The name of the database to be created is *dbName*.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCIN0053E An error occurred when creating the database *dbName*. Review the log files for more information.

Explanation

The installation program cannot recover from the installation error.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: TPC_installation_directory\logs\
- For UNIX or Linux operating system: TPC_installation_directory/logs/

TPC_installation_directory is where IBM Spectrum Control is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

Correct the problem. If you cannot correct the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0055E Db2 is not installed or the Db2 profile was not sourced before installing IBM Spectrum Control.

Explanation

Db2 must be installed before installing IBM Spectrum Control. For AIX or Linux systems, you must source the Db2 profile before installing IBM Spectrum Control.

Action

If Db2 is not installed, install Db2. If you are running on AIX or Linux, go to the IBM Knowledge Center and search on the text string "Starting Installation". Follow the instructions on sourcing the Db2 profile then run the installation program again.

BPCIN0056E The user name *userID* does not have write permission on the default database path configuration parameter **DFTDBPATH:** *dftdbpath*.

Explanation

An error occurred because the specified user name does not have write permission on the default database path configuration parameter. For the installation, IBM Spectrum Control must have write permission on this parameter.

Action

Set the write permission for the specified user name on DFTDBPATH {1}. For more information about how to set this parameter, go to the IBM Knowledge Center and search on the parameter.

BPCIN0057E The user name *userID* is not in an operating system group that has Db2 SYSADM authority. Before you run your DB2 installation, validate that the Windows Server Hostname is 15 characters or less in length.

Explanation

An error occurred because the specified user name is not in an operating system group that has Db2 SYSADM authority.

The problem also could be due to the Windows Server Hostname more than 15 characters in length.

Action

Add the specified user name to an operating system group that has Db2 SYSADM authority.

Before you run your DB2 installation, validate that the Windows Server Hostname is 15 characters or less in length.

BPCIN0059W There are long file names in the IBM Spectrum Control installation images.

Explanation

The reason for the long file names could be a result of using an incorrect version of the AIX tar program. Use the GNU tar program to extract files. The AIX tar program might truncate long file names, which can cause errors in the IBM Spectrum Control installation.

Action

Use the GNU tar program, instead of the native AIX tar program, to extract files from the IBM Spectrum Control installation images.

BPCIN0060E An error occurred during the uninstallation of component. Review the log files for more information.

Explanation

The uninstallation program cannot recover from the uninstallation error.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: %HOMEPATH% directory
- For UNIX or Linux operating systems: \$HOME directory

The log files are:

- msgTPCStdout.txt
- msgTPCStderr.txt

BPCIN0061E An invalid host name or IP address was specified for the Data server.

Explanation

The installation program requires a fully qualified host name or valid IP address for the Data server.

An example of a fully qualified host name is: `myuser.ibm.com`.

An example of a valid IP address is: `127.0.0.1`.

Action

Enter a fully qualified host name or valid IP address for the Data server.

BPCIN0062E The IPv6 internet protocol is not enabled on the specified host computer.

Explanation

To use the IPv6 internet protocol for a host computer, you must have this support enabled.

Action

Enable the IPv6 internet protocol on the specified host computer and run the installation program again.

BPCIN0063E The Data server is not running. For more information, go to the IBM Knowledge Center and search on the message code.

Explanation

The installation program cannot continue because the Data server is not running.

Action

Verify that the Data server is running:

- For Windows operating systems, open the Windows Services window and see if the following service is running: IBM Spectrum Control - Data Server.
- For UNIX operating systems, check to see if the TSRM process is running.
- To verify that the TSRM process is running, run the following command: `ps -ef|grep TSRM`.

If the Data server is not running, start the Data server and run the installation program again.

To start the Data server on Windows, run this command `TPC_installation_directory\scripts\startTPCData.bat`.

To start the Data server on UNIX, run this command `TPC_installation_directory/scripts/startTPCData.sh`.

`TPC_installation_directory` is where IBM Spectrum Control is installed.

BPCIN0064E The Data Server is not running at the specified host address or port: host *host_address*, port *port*. For more information, go to the IBM Knowledge Center and search for the message code.

Explanation

The Data Server is not running at the specified host or port or the address details are incorrect.

Action

Verify the host address details. If they are correct, verify that the Data server is running:

- For Windows operating systems, open the Windows Services window and see if the following service is running: IBM Spectrum Control - Data Server.
- For UNIX operating systems, check to see if the TSRM process is running.
- To verify that the TSRM process is running, run the following command: `ps -ef|grep TSRM`.

If the Data server is not running, start the Data server and run the installation program again.

To start the Data server on Windows, run this command `TPC_installation_directory\scripts\startTPCData.bat`.

To start the Data server on UNIX, run this command `TPC_installation_directory/scripts/startTPCData.sh`.

`TPC_installation_directory` is where IBM Spectrum Control is installed.

BPCIN0066E Errors occurred during the installation of the IBM Spectrum Control GUI. Review the log files for more information.

Explanation

The connection with WebSphere Application Server Liberty failed.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: `TPC_installation_directory\logs\`
- For UNIX or Linux operating system: `TPC_installation_directory/logs/`

`TPC_installation_directory` is where IBM Spectrum Control is installed.

The log files are:

- `msgTPCInstall.log`
- `traceTPCInstall.log`

Correct the problem and run the installation program again. If you cannot correct the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0068E Errors occurred during the uninstallation of the IBM Spectrum Control GUI. Review the log files for more information.

Explanation

Errors occurred when the uninstallation program tried to delete the WebSphere Application Server Liberty data source object.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: %HOMEPATH% directory
- For UNIX or Linux operating systems: \$HOME directory

The log files are:

- msgTPCStdout.txt
- msgTPCStderr.txt

Correct the problem and run the installation program again. If you cannot correct the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0069E Errors occurred during the configuration of Tivoli Common Reporting for IBM Spectrum Control. Review the log files for more information.

Explanation

IBM Spectrum Control could not configure the database for Tivoli Common Reporting.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: TPC_installation_directory\logs\
- For UNIX or Linux operating system: TPC_installation_directory/logs/

TPC_installation_directory is where IBM Spectrum Control is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

Correct the problem and run the installation program again. If you cannot correct the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0070E Errors occurred during the configuration of the IBM Spectrum Control data model in Tivoli Common Reporting. Review the log files for more information.

Explanation

Errors occurred when the installation program tried to import the IBM Spectrum Control data model to Tivoli Common Reporting.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: TPC_installation_directory\logs\
- For UNIX or Linux operating system: TPC_installation_directory/logs/

TPC_installation_directory is where IBM Spectrum Control is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

Correct the problem and run the installation program again. If you cannot correct the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0071E An error occurred because the installation program could not find the Db2 DFTDBPATH variable. For more information, go to the IBM Knowledge Center and search for the message code.

Explanation

DFTDBPATH is a Db2 variable that contains the default file path used to create the databases managed by the database manager. The installation program cannot recover from this error.

Action

Contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0072E The validation for user name *userID* has failed. Check to see if this user name exists or if Db2 is running.

Explanation

The validation for the specified user name has failed because the user name might not exist on the system or Db2 is not running.

Action

Check to see if the user name exists on the system or if Db2 is running. Correct the problem and run the installation program again.

BPCIN0074E The port range validation failed because the port value is not numeric.

Explanation

The port range validation failed because the port value is not numeric.

Action

Enter a numeric value for the port.

BPCIN0075E An invalid GUID 0xFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF was found. Update or uninstall the GUID.

Explanation

An invalid GUID 0xFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF was found.

Action

Update or uninstall the GUID and run the installation program again.

BPCIN0076E IBM Spectrum Control could not read the GUID. See the GUID installation log for an explanation of the error.

Explanation

IBM Spectrum Control could not read the GUID. See the GUID installation log for an explanation of the error.

Action

Review the guidInstallOutput.log file for an explanation of the error. The log file is located in the following default directory:

- For Windows operating systems: %temp% directory
- For UNIX or Linux operating system: /tmp directory

Correct the problem and run the installation program again. If you cannot correct the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0077E The name specified for the database is not valid because it contains a space or is blank.

Explanation

The database name that was specified is not valid because it contains a space or is blank.

Action

Enter a valid database name.

BPCIN0078E The database name *database_name* is not valid. The name can only contain the following characters: a-z, A-Z or 0-9.

Explanation

The database name can have the following characters:

- a-z
- A-Z
- 0-9

Action

Enter a valid database name.

BPCIN0079E The first character or characters in the database name are not valid. The name must not begin with a number or the letters: SYS, DBM, or IBM.

Explanation

The first character or characters in the database name must not begin with:

- A number
- The letters: SYS, DBM, or IBM

The valid characters for the database name are:

- a-z
- A-Z
- @, #, or \$

Action

Enter a valid database name.

BPCIN0080E The database name specified is not valid because an existing database has the same name: *database_name*.

Explanation

You cannot specify a database name that has the same name as an existing database.

Action

Enter another database name.

BPCIN0081E The database name *database_name* is too long. The name can be 1 - 8 characters in length with no spaces.

Explanation

The database name that was specified is too long. The value can be 1 - 8 characters in length with no spaces.

Action

Enter a valid database name.

BPCIN0082E The specified host name or IP address is not a remote host. Enter a remote host name or IP address.

Explanation

The specified host name or IP address must be a remote host name or IP address.

Action

Enter a remote host name or IP address.

BPCIN0083E The server could not connect to the remote database. Verify that the port is correct and that the database is running on the remote server. Also verify that the internet protocol connection between the server and remote database is compatible.

Explanation

The server could not connect to the remote database. The reason for this issue can be:

- The port specified does not match the port in the Db2 configuration file for the remote database
- The database is not running on the remote server
- The internet protocol connection between the server and remote database is not compatible

Action

Verify that the remote database port is correct and that the database is started. Run the installation program again. To verify the database port, run the following command from a command window: db2 get dbm cfg. This command opens the Db2 database manager configuration file. Search for "SVCENAME." Open the services file in this directory:

- For the Windows operating system, see %WINDOWS%\system32\drivers\etc\services.
- For the UNIX or Linux operating system, see /etc/services.

For information about valid internet protocol configurations, go to the IBM Knowledge Center and search for "Valid configurations for IPv4 and IPv6 systems."

BPCIN0084E The database name is not valid because there is no database on the specified server with this name.

Explanation

You must specify a database that exists.

Action

Check the remote system for a valid database.

BPCIN0085E The installation program was unable to connect to the remote database. Check the remote server to verify that you can connect to Db2. For more information about Db2 connection issues, go to the IBM Knowledge Center at <http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome> .

Explanation

An error occurred when the installation program tried to create a JDBC connection to the remote database. A reason for this error might be that the database was not started on the remote server.

Action

Check the remote server and verify that you can connect to Db2. For more information about Db2 connection issues, go to the IBM Knowledge Center at <http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome> and search for "Troubleshooting Db2 Server."

BPCIN0086E The connection to the remote database failed because the user name or password is not valid.

Explanation

You must specify a valid user name and password to connect to the remote database.

Action

Check your remote database for a valid user name and password.

BPCIN0087E The remote host name or IP address is not valid.

Explanation

You must specify a valid fully qualified host name or IP address.

Action

Check your remote host for the host name or IP address.

BPCIN0088E The database repository was not found on the remote server.

Explanation

The database repository must be installed on the remote server.

Action

Install the database repository on the remote server and run the installation program again.

BPCIN0089E The version "*retrivedVersion*" of the remote database repository is not at the correct level. Upgrade the remote database repository to version "*requiredVersion*", then upgrade the remaining components.

Explanation

The remote database repository is not at the correct version. Upgrade the remote database repository to the correct version, then upgrade the remaining components.

Action

Upgrade the remote database repository to the correct version, then upgrade the remaining components.

Install the correct version of the database repository on the remote server and run the installation program again.

BPCIN0090E The name specified for the database is not valid because it contains a space.

Explanation

The database name that was specified is not valid because it contains a space.

Action

Enter a valid database name.

BPCIN0091E The file path *file_path* specified is not an absolute path or is not a directory or the partition does not exists.

Explanation

The file path specified must be an absolute path and not a relative path.

An example of an absolute path for Windows is: C:\Program Files\IBM\TPC.

An example of an absolute path for UNIX or Linux is: /opt/IBM/TPC.

Action

Enter a valid value for the path.

BPCIN0092E The first directory *directory_name* specified is not valid.

Explanation

The first directory in the path has a name in the form SQLNNNNN, where NNNNN is a value from 00001 to 99999.

Action

Enter a valid value for the path.

BPCIN0093E The path string *path_string* is too long. The path string cannot be longer than 242 bytes.

Explanation

The path string is longer than 242 bytes.

Action

Enter a valid value for the path.

BPCIN0094E The user name *user_name* does not have write permission on the specified database path: *database_path*.

Explanation

The specified user name does not have write permission on the specified database path.

Action

Set the write permission for the specified user name on the database path. For information about how to set this parameter, go to the IBM Knowledge Center at <http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome> and search for the db2extsec command.

BPCIN0095E The user name *user_name* does not have write permission on the specified log location: *log_location*.

Explanation

The specified user name does not have write permission on the specified log location.

Action

Set the write permission for the specified user name for the log location. For information about how to set this parameter, go to the IBM Knowledge Center at <http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome> and search for the db2extsec command.

BPCIN0096E The database path cannot be blank. Enter a valid database path.

Explanation

You must specify a database path.

Action

Enter a valid database path.

BPCIN0097E The log location cannot be blank. Enter a valid log location.

Explanation

You must specify a log location.

Action

Enter a valid log location.

BPCIN0098E Enter 10 or fewer paths.

Explanation

The maximum number of allowed database locations has been reached

Action

Enter 10 or fewer paths.

BPCIN0099E You have entered *pathNumber* directories. A maximum of 10 directories can be specified.

Explanation

The maximum number of allowed database locations has been exceeded

Action

A maximum of 10 directories can be specified.

BPCIN0100W The database log files and database are in the same location. Click Yes to change the database log files or database location. Click No to ignore this message.

Explanation

To ensure that the log files are on a physical disk that does not have high I/O activity, change the database log files or database location. For instance, avoid placing the logs on the same disk as the operating system or high volume databases. In this way, data is logged more efficiently because the logging process does not have to wait for other I/O to complete.

Action

Change the database log path or database location.

BPCIN0101E The location "*location*" that was specified contains Db2 logs or log files.

Explanation

The directory that was specified contains Db2 logs or log files. The directory must not contain Db2 logs or log files. Examples of Db2 logs or log files are: S0000000.LOG, S0000001.LOG, or SQLLPATH.TAG.

Action

Select a directory that does not contain Db2 logs or log files.

BPCIN0102E The Db2 profile for the "*db2_instance*" instance was not run before installing IBM Spectrum Control.

Explanation

The Db2 profile must be run before you install IBM Spectrum Control with the specified Db2 instance.

Action

For UNIX and AIX systems, run the Db2 profile with the specified Db2 instance before you install IBM Spectrum Control.

BPCIN0103E The user name *userID* is not in the system group *adminGroup*.

Explanation

The user name to install IBM Spectrum Control must be in the system group.

Action

Add the user name to the system group and continue with the installation.

BPCIN0104I Resuming a failed install and installing the remaining components.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCIN0105E A reboot must be done before continuing with the installation.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCIN0106E There are invalid header files in the installation images. You may might be using the AIX tar program instead of the GNU tar program to extract files from the installation images. For more information, see IBM Knowledge Center and search by the error message code.

Explanation

The AIX tar program can result in invalid header files, because this tar program places the header files in the installation images. This can cause errors in the installation.

Action

Use the GNU tar program rather than the native AIX tar program to extract files from the installation images.

Download, install, and use the GNU tar program, version 1.14 or later. The GNU tar program can be downloaded from <http://www-03.ibm.com/systems/power/software/aix/linux/toolbox/alpha.html>. The GNU tar program must be installed as the default tar utility in the installation path. The default installation location for this tar program is /usr/local/bin.

BPCIN0107E The database path exists in the database path list. You must specify a unique database path.

Explanation

You must specify a unique database path and cannot reuse an existing path.

Action

Enter a valid database path.

BPCIN0108E An error occurred during the upgrade of the *component*. Review the log files in the following directory for an explanation of the error: *location*.

Explanation

The specified component was not upgraded because of errors. One possible cause of these errors is the Windows Administrator user that is installing the upgrade does not have Windows Debug Programs rights.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: TPC_installation_directory\logs\
- For UNIX or Linux operating systems: TPC_installation_directory/logs/

TPC_installation_directory is the directory where the product is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

To verify that the Windows Administrator user has Debug Programs rights, go to Administrative Tools > Local Security Policy > Local Policies > User Rights Assignment and right click Debug programs in the Policy list. If necessary, add the administrator user or group. Log off Windows and log back in again.

Correct any other problems identified in the log files and run the installation or upgrade program again. If you cannot resolve the problem, contact IBM Software Support.

Related reference

- [Getting support](#)
- [Enabling and disabling legacy protocol \(SSLv3 and MD5 hash\)](#)

BPCIN0109E An unexpected error occurred. IBM Spectrum Control cannot resolve this error. For more information, review the log files and go to the IBM Knowledge Center.

Explanation

An error occurred during the installation program.

Action

Review the log files in one of the following directories:

- For Windows operating systems: TPC_installation_directory\logs\
- For AIX or Linux operating systems: TPC_installation_directory/logs/

TPC_installation_directory is where IBM Spectrum Control is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

Resolve the issue and run the installation program again. If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0112I Upgrading the components: *upgrade_components*.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCIN0116I The system is upgrading the *component* component.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCIN0117I The system completed the upgrade of the *component* component.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCIN0120E The validation of user name and password could not be completed

Explanation

An error occurred when validating the user name and password

Action

Verify that the user name and password matches the operating system user name and password. Continue with the installation.

BPCIN0121E The installation type entered is not supported.

Explanation

The installation type entered is not supported.

Action

Enter a valid installation type. Continue with the installation.

BPCIN0122E The `TPCCommon.dll` or `libTPCCommon.so` library could not be found or loaded.

Explanation

The `TPCCommon.dll` or `libTPCCommon.so` library could not be located or loaded .

Action

Check the IBM Spectrum Control installation image for the missing libraries and make sure that they are in the installation image. Continue with the installation.

BPCIN0123E The user has no administrative rights to install or uninstall IBM Spectrum Control.

Explanation

A user must have administrative rights to install or uninstall IBM Spectrum Control.

Action

Log in with a valid user name that has the correct administrative rights to install or uninstall IBM Spectrum Control.

BPCIN0124E The operating system cannot find the Db2 Service name.

Explanation

The database instance information cannot be found.

Action

Check the Db2 instance configuration.

BPCIN0125E The installation program cannot find the Db2 instance name. Check the Db2 instance configuration.

Explanation

The name of the Db2 instance cannot be found.

Action

Check the Db2 instance configuration.

BPCIN0126E Unable to find Db2 installation path. Check your Db2 configuration.

Explanation

The Db2 installation path cannot be found on your system.

Action

Check your Db2 configuration.

BPCIN0127E The service name configuration parameter cannot be queried from the Db2 Database Manager.

Explanation

The service name configuration parameter cannot be queried from the Db2 Database Manager.

Action

Check your Db2 instance configuration.

BPCIN0129E The version that is installed cannot be upgraded because it does not meet the minimum build version for the database repository. The minimum build version supported is *build_version*.

Explanation

The upgrade failed when you attempted to upgrade to a different build within the same version of software. The installed version does not meet the minimum build requirements for the database repository.

Action

If the database does not meet the minimum build requirements, check the version.txt file. The version.txt file displays the version that you have installed on your system. To display the version.txt file, go to the following directory:

- For Windows operating systems: TPC_installation_directory
- For AIX or Linux operating system: TPC_installation_directory

TPC_installation_directory is where IBM Spectrum Control is installed.

Upgrade the software that is installed to the minimum build version or later before upgrading to the current build.

BPCIN0130E The database administrator password in IBM Spectrum Control does not match the database administrator password for Db2. Run the changepasswords tool to update the database administrator password.

Explanation

The database administrator password in IBM Spectrum Control must match the database administrator password for Db2.

Action

Run the changepasswords tool to change the password for the database administrator user.

In a command prompt, go to the following directory:

- For Windows operating systems: TPC_installation_directory\service\
- For AIX or Linux operating system: TPC_installation_directory/service/

TPC_installation_directory is where IBM Spectrum Control is installed.

Run this command:

- For Windows operating systems: changepasswords.bat
- For AIX or Linux operating system: changepasswords.sh

Run the installation program again.

For information about the changepasswords tool, go to the IBM Knowledge Center and search for "using password tool".

BPCIN0131E An invalid host name or IP address was specified for the Device Server.

Explanation

The installation program requires a fully qualified host name or valid IP address for the Device server.

An example of a fully qualified host name is: myuser.example.mycompany.com.

An example of a valid IP address is: 127.0.0.1.

Action

Enter a fully qualified host name or valid IP address for the Device server.

BPCIN0132E The Device Server is not running at the specified host address or port: host *host_address*, port *port*.

Explanation

The Device server is not running at the specified host or port or the address details are incorrect.

Action

Verify the host address details. If they are correct, verify that the Device server is running:

- For Windows operating systems, open the Windows Services window and see if the following service is running: IBM Spectrum Control - Device Server.
- For AIX or Linux operating systems, run this command: `ps -ef|grep deviceServer`. If the server is running, you will see a list of processes.

If the Device server is not running, start the Device server and run the installation program again.

To start the Device server on Windows, run this command `TPC_installation_directory\scripts\startTPCDevice.bat`.

To start the Device server on AIX and Linux, run this command `TPC_installation_directory/scripts/startTPCDevice.sh`.

TPC_installation_directory is where IBM Spectrum Control is installed.

BPCIN0134E The installation program could not stop the . Please manually stop the .

Explanation

The installation program could not stop the server.

Action

Try to manually stop the server and continue with the installation or upgrade.

BPCIN0135E Tivoli Storage Productivity Center version is installed. Before you can upgrade to version 5, you must upgrade Tivoli Storage Productivity Center to version 4.

Explanation

You cannot upgrade Tivoli Storage Productivity Center version 3 directly to version 5. You must first upgrade version 3 to version 4, then upgrade version 4 to version 5.

Action

For information about upgrading to version 4, see the Tivoli Storage Productivity Center Knowledge center version 4.2.2. Search for Upgrading and migrating the Tivoli Storage Productivity Center family. After you upgrade to version 4, upgrade to version 5.

BPCIN0136E The version "*retrivedVersion*" of the remote database repository is not at the correct build level. Upgrade the remote database repository to the correct build level, then upgrade the remaining version "*requiredVersion*" components.

Explanation

The remote database repository is not at the correct build level. Upgrade the remote database repository to the correct build level, then upgrade the remaining components.

Action

Upgrade the remote database repository to the correct build level, then upgrade the remaining components.

BPCIN0137E The version "*retrivedVersion*" of the remote database repository is not at the correct build level. Ensure that the installation image for the remaining version "*requiredVersion*" components is at the same build level as the remote database repository, then upgrade the remaining components.

Explanation

The remote database repository is not at the correct build level. Ensure that the installation image for the remaining components is at the same build level as the remote database repository, then upgrade the remaining components.

Action

Ensure that the installation image for the remaining components is at the same build level as the remote database repository, then upgrade the remaining components.

BPCIN0138E The common user name password does not match the password for the existing user name *user_name* on the system. Run the `changepasswords` tool to change the common user name password for IBM Spectrum Control.

Explanation

Before you can run the installation program, the password specified must match the password for the user name that exists on the system.

Action

Verify that the password specified and the user name password on the system matches. If the passwords do not match, run the `changepasswords` tool.

Run the `changepasswords` tool to change the password for the common user name.

In a command prompt, go to the following directory:

- For Windows operating systems: `TPC_installation_directory\service\`
- For AIX or Linux operating system: `TPC_installation_directory/service/`

Where `TPC_installation_directory` is where IBM Spectrum Control is installed.

Run this command:

- For Windows operating systems: `changepasswords.bat`
- For AIX or Linux operating system: `changepasswords.sh`

Run the installation program again.

For information about the `changepasswords` tool, go to the IBM Knowledge center and search for "Using the password tool".

BPCIN0139I The User Migration Tool completed successfully.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCIN0140E The User Migration Tool could not be started. Start the User Migration Tool from the graphical user interface after the upgrade is finished.

Explanation

The upgrade will continue. You can start the User Migration Tool after the upgrade completes.

Action

No action is required.

BPCIN0143E Errors occurred during the upgrade of the GUI. Review the log files for more information.

Explanation

The upgrade program could not update the TPC-GUI.war file.

Action

Review the log files in one of the following directories:

- For Windows operating systems: TPC_installation_directory\logs\
- For AIX or Linux operating systems: TPC_installation_directory/logs/

TPC_installation_directory is where the product is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

Correct the problem and run the upgrade program again. If you cannot correct the problem, contact IBM Software Support.

Related reference

-  [Getting support](#)

BPCIN0146E The installation language specified is not a supported language.

Explanation

You must specify a language that is supported.

Action

Specify one of the following languages:

- cs - Czech
- de - German
- en - English
- es - Spanish
- fr - French
- it - Italian
- ja - Japanese
- ko - Korean
- pl - Polish
- pt_BR - Brazilian Portuguese
- hu - Hungarian
- ru - Russian
- zh_CN - Simplified Chinese
- zh_TW - Traditional Chinese

BPCIN0148E IBM Spectrum Control cannot be installed, because the physical memory size on this computer is too small. The minimum memory size is *minMemoryProduction*. Increase the physical memory on this computer and run the installation program again.

Explanation

The minimum memory size requirements for IBM Spectrum Control have not been met.

Action

Increase the physical memory on this computer and run the installation program again.

BPCIN0149W The physical memory size on this computer is below the minimum requirements that are specified for a production system. The minimum memory size for a production system is *minMemoryProduction*. You must increase the physical memory on this computer and run the installation program again. If you install IBM Spectrum Control with lower memory, you can only use it in an evaluation environment.

Explanation

The suggested memory size requirements for IBM Spectrum Control are not met.

Action

Increase the physical memory on this computer and run the installation program again. If you install IBM Spectrum Control with lower memory, you can only use it in an evaluation environment.

BPCIN0150E The installation program cannot determine the amount of physical memory on the system.

Explanation

The installation program cannot determine the amount of physical memory on the system for this installation.

Action

Contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0151E You tried to install IBM Spectrum Control on an unsupported operating system. For more information on the supported operating systems, go to following link.

Explanation

IBM Spectrum Control cannot be installed on this operating system because the version of the operating system is not supported.

Action

For information about the operating systems that are supported, go to the IBM Knowledge Center and search for "software requirements for operating systems."

Related reference

- [IBM Spectrum Control 5.3.x - Servers, Agents, and Browsers](#)

BPCIN0152E IBM Spectrum Control cannot determine if the operating system is supported.

Explanation

An error occurred while verifying that the operating system is supported.

Action

If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0153E The installation program could not rename the jre folder. To continue the installation, stop all Java processes that access the jre directory: *jre_folder*.

Explanation

An error occurred while renaming JRE .

Action

Please ensure that all java processes accessing the jre directory are stopped.

BPCIN0159E The web server data source creation failed.

Explanation

An error occurred while creating web server data source.

Action

For more information about the error, review the log files in one of the following directories:

- For Windows operating systems: TPC_installation_directory\logs\
- For AIX or Linux operating systems: TPC_installation_directory/logs/

TPC_installation_directory is where IBM Spectrum Control is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

After you resolve the issue, run the installation program again. If you cannot resolve the issue, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0163E The service name configuration parameter cannot be queried from the Db2 Database Manager. Reboot the machine after Db2 installation. SQL Error Message is *sqlerrmsg*

Explanation

The service name configuration parameter cannot be queried from Db2 Database Manager.

Action

Check your Db2 instance configuration.

BPCIN0164E The service name configuration parameter cannot be queried from the Db2 Database Manager. SQL Error Message is *sqlerrmsg*

Explanation

The service name configuration parameter cannot be queried from Db2 Database Manager.

Action

Check your Db2 instance configuration.

BPCIN0165E The password for user name has expired. You must change the password, or select another user name, to install IBM Spectrum Control.

Explanation

The password for the entered user name has expired. Please ensure that the password of the entered user account is not expired.

Action

Change password for corresponding user name.

BPCIN0166W The password for user name cannot be checked for expiration. Please ensure that it is not expired.

Explanation

The password cannot be checked for expiration. Please ensure manually that it is not expired

Action

Expired password will cause an error during installation.

BPCIN0167E You cannot upgrade the license to the same or lower level.

Explanation

If the license key file provided is the same or lower than the installed IBM Spectrum Control license type, you cannot upgrade your license.

Action

Specify a higher license level.

BPCIN0168E An error occurred while checking the user name and password. Review the log files for more information.

Explanation

An error occurred while checking the user name and password.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: TPC_installation_directory\logs\
- For AIX or Linux operating systems: TPC_installation_directory/logs/

TPC_installation_directory is where IBM Spectrum Control is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

Correct the problem and run the installation program again. If you cannot resolve the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0169E The location "*location*" that should be specified for the license is not correct.

Explanation

Specify a valid location for the license file.

Action

Select a license file.

BPCIN0170E IBM Spectrum Control cannot upgrade the license because the upgrade must be done on the server system.

Explanation

To upgrade your license, you must upgrade the license on the IBM Spectrum Control servers system.

Action

Before you upgrade the license, ensure that you have the IBM Spectrum Control servers installed on your system.

BPCIN0173E An error occurred during the deployment of the *file_name* file.

Explanation

The IBM Spectrum Control installation program could not deploy the war file specified in the message text.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: TPC_installation_directory\logs\
- For AIX or Linux operating systems: TPC_installation_directory/logs/

TPC_installation_directory is where IBM Spectrum Control is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

After you resolve the issue, you must run the installation program again. If you cannot resolve this issue, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0176E An error occurred during Db2 catalog creation, and the catalog was not created.

Explanation

The catalog entry either exists or the Db2 command is not available.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: TPC_installation_directory\logs\
- For AIX or Linux operating systems: TPC_installation_directory/logs/

TPC_installation_directory is where IBM Spectrum Control is installed.

The log files are:

- msgTPCInstall.log
- traceTPCInstall.log

If the log file indicates that the catalog entry exists, remove the catalog entry. To remove the catalog entry, run the following commands:

- db2 uncatalog node node_name
- db2 uncatalog db database_alias_name
- db2 terminate

BPCIN0177E You cannot upgrade IBM Spectrum Control with a license key file that is at a lower level than the installed license key file.

Explanation

If the license key file you are using to upgrade is a lower level than the installed key file, you cannot upgrade.

Action

To upgrade, you must use a license key file with a higher level.

BPCIN0178E The provided license key *file_name* is invalid.

Explanation

The provided license key is invalid.

Action

Provide a valid license file.

BPCIN0179E An error occurred during Db2 catalog deletion, and the catalog was not removed.

Explanation

The catalog entry exists, but it could not be deleted or the Db2 command is not available.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: %HOMEPATH% directory
- For UNIX or Linux operating systems: \$HOME directory

TPC_installation_directory is where IBM Spectrum Control is installed.

The log files are:

- msgTPCStdout.txt
- msgTPCStderr.txt

BPCIN0181E The web server data source testing has failed.

Explanation

An error occurred during the web server data source testing.

Action

An error occurred during the web server data source testing. If the problem continues, contact IBM Support.

Related reference

-  [Getting support](#)

BPCIN0182E TPC-GUI.war data source testing failed.

Explanation

An error occurred while testing the TPC-GUI.war data source.

Action

Review the installation log files.

BPCIN0184E The remote database repository already contains data. Install a new IBM Spectrum Control remote database repository first, then install the remaining IBM Spectrum Control components.

Explanation

The remote database repository must not contain data associated with another installation.

Action

Install a new IBM Spectrum remote database repository first, then install the remaining IBM Spectrum Control components.

BPCIN0185E The remote database repository for version *oldVersion* is incompatible with IBM Spectrum Control Version *requiredVersion*. Upgrade remote database repository to Version *requiredVersion* first, and then install the remaining IBM Spectrum Control components.

Explanation

The remote database repository is incompatible with IBM Spectrum Control.

Action

Install a IBM Spectrum Control remote database repository first, then install the remaining IBM Spectrum Control components.

BPCIN0190E You cannot use non-standard characters, such as a space between characters or an underscore, in a host name. You must

enter a host name with standard characters and try again.

Explanation

You cannot use non-standard characters, such as a space between characters or an underscore, in a host name.

Action

You must enter a host name with standard characters and try again.

BPCIN0191E You cannot install Db2 in a directory that starts with the letter a.

Explanation

If the directory in which you want to install Db2 starts with the letter a, for example, C:\abc\IBM\ABCD, the installation will fail.

Action

Select a directory that does not begin with the letter a, such as C:\ABC\IBM\ABCD.

BPCIN0193E The version "*retrivedVersion*" of the remote database repository is not at the correct level. Install a new Version "*requiredVersion*" remote database repository at the correct level, then install the remaining IBM Spectrum Control components.

Explanation

The remote database repository is not at the correct level. Install a new IBM Spectrum Control remote database repository at the correct level, then install the remaining IBM Spectrum Control components.

Action

Install a new IBM Spectrum Control remote database repository at the correct level, then install the remaining IBM Spectrum Control components.

BPCIN0195E The version "*retrivedVersion*" of the database repository *dbHost* is not at the correct build level. Install a new Version "*requiredVersion*" database repository at the correct build level, then install the remaining IBM Spectrum Control components.

Explanation

The database repository for IBM Spectrum Control is not at the correct build level. Install a new IBM Spectrum Control database repository at the correct build level, then install the remaining IBM Spectrum Control components.

Action

Install a new database repository at the correct build level, then install the remaining IBM Spectrum Control components.

BPCIN0198E The path to installation image "*installDirectory*" contains the invalid character "*unallowedCharater*". You must change the directory name so that it contains valid characters.

Explanation

The path to the installation image contains invalid characters.

Action

You must select a valid installation directory.

BPCIN0199E The installation program cannot validate the host name because the fully qualified domain name (FQDN) cannot be retrieved for the host.

Explanation

The host name cannot be validated because the local fully qualified domain name cannot be retrieved for the host. Either the host name is incorrect or the DNS is providing an invalid configuration.

Action

Review and correct the Domain Name System (DNS) configuration on the host system or specify a valid host name. Continue with the installation.

BPCIN0200E The installation program does not allow host name specified as IP address. Please ensure that a fully qualified domain name (FQDN) is provided.

Explanation

The installation program does not allow host name specified as IP address. That's because SSO and LIC function will not work, if the system is not configured with FQDN.

Action

Review and correct the Domain Name System (DNS) configuration on the host system or specify a valid host name. Continue with the installation.

BPCIN0202E Runtime errors have occurred during the IBM Spectrum Control preinstallation process. The installation program cannot recover from this error.

Explanation

The installation program cannot recover from this error.

Action

Contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0203E The installation program does not support backslash character in the user name. In case it is a Windows domain account, please specify just the user name without using the "Domain_name\\" prefix.

Explanation

The installation program does not support backslash character in the user name. In case it is a Windows domain account, please specify just the user name without using the "Domain_name\\" prefix.

Action

Enter a valid user name and continue with the installation.

BPCIN0205E The IBM Spectrum Control installation program could not find the directory "*installDirectory*" on the installation image. For more information, go to the IBM Knowledge Center and search on the message code.

Explanation

The installation program could not find the directory or directories on the installation image for one of the following reasons:

- An error occurred during the installation image download.
- A directory was deleted from the installation image.
- An error occurred with a network drive.

Action

If you downloaded an installation image, verify that this image was downloaded correctly.

Correct the problem and restart the installation.

If the problem persists, contact IBM Software Support.

Related reference

-  [Getting support](#)

BPCIN0206E The IBM Spectrum Control installation program was unable to retrieve a fully qualified domain name (FQDN) for the host. You must configure the host system with an FQDN.

Explanation

The host name could not be validated because the local FQDN could not be retrieved for the host. The host name is either incorrect or the Domain Name System (DNS) is providing an invalid configuration.

Action

Correct the DNS configuration on the host system or specify a valid host name and continue with the installation.

BPCIN0207E You cannot upgrade when the stand-alone GUI is running. Stop the stand-alone GUI and click OK to continue the upgrade or click Quit to exit the installation program.

Explanation

You cannot upgrade when the stand-alone GUI is running. To continue the upgrade, stop the stand-alone GUI .

Action

Stop the stand-alone GUI and click OK to continue the upgrade or click Quit to exit the installation program.

BPCIN0208E The fully qualified domain name(FQDN) retrieved for the host contains non-standard characters, such as a space between characters or an underscore. You must configure the host system with an FQDN that contains standard characters.

Explanation

You cannot use non-standard characters, such as a space between characters or an underscore, in a host name.

Action

Configure the host system with an FQDN that contains standard characters and continue with the installation.

BPCIN0209E The installation program does not allow a user that is present in both the Windows Domain and the local Operating System repositories. Install with a user name that is only present in one of these repositories.

Explanation

The installation program does not allow a user name that present in both the Windows Domain (or Active Directory) and local Operating System repositories.

Action

Ensure that the specified user is present in either the Windows Domain or the local Operating System repositories.

BPCIN0210E The Db2 "*sourcedDB2profile*" profile has not been loaded in the .profile file for user *commonUser*. This profile must be loaded before you install IBM Spectrum Control.

Explanation

The Db2 profile must be loaded before you install IBM Spectrum Control with the specified Db2 instance.

Action

For UNIX and AIX operating systems, before you install IBM Spectrum Control, load the Db2 profile in the .profile file of the user name employed to install the IBM Spectrum Control database component.

BPCIN0211E The technology level or the service pack level of this operating system is not supported and must be upgraded to a supported level. The detected operating system version is *os_version*

Explanation

IBM Spectrum Control cannot be installed on this system because the technology level or the service pack level of this operating system is not supported. The technology level or the service pack level must be upgraded to a supported level.

Action

For information about supported operating systems, go to the IBM Knowledge Center and search for "software requirements for operating systems."

BPCIN0214E Cannot complete a fresh install without a valid license.

Explanation

Cannot complete a fresh install without a valid license.

Action

You must download an install image with a license.

BPCIN0215E The path *extractorDirectory* where IBM Spectrum Control is extracted has *extractorDirLength* characters and the maximum number of characters allowed is 260. You must shorten this path.

Explanation

The Windows operating system imposes a maximum length of 260 characters for a file path.

Action

You must shorten the path to the directory where IBM Spectrum Control is extracted.

BPCIN0217E The Windows registry check indicates that .NET 3.5 or higher is not available on this Windows server. You must install .NET 3.5 or higher to continue. To install .NET 3.5 type the following commands in a 64 bit windows powershell: `Import-Module ServerManager` `Add-WindowsFeature as-net-framework`

Explanation

You must install .NET 3.5 to continue.

Action

To install .NET 3.5, type the following commands:

BPCIN0219E The domain configuration is invalid. To resolve the issue, complete the following steps: Disable the Windows Firewall service. Start or restart the Computer Browser service on this domain member computer and on the domain controller computer. If the service has a Stopped or Disabled status on the domain controller computer, you must restart Computer Browser service on the domain member computer after you start the service on the domain controller computer. In a Windows command window, run the `net view` command and verify that there are no errors. Reinstall IBM Spectrum Control.

Explanation

The domain configuration is invalid.

Action

The domain configuration is invalid. To resolve the issue, complete the following steps:

1. Disable the Windows Firewall.
2. Start the Computer Browser service.
3. In a command window, run the "net view" command and verify that there are no errors in the output.
4. Reinstall IBM Spectrum Control.

BPCIN0220E The current login user *loginuser* is not an administrator or a member of the domain administrator group. The installation program cannot start.

Explanation

This user name is not an administrator or a member of the domain administrator group.

Action

To start the installation program, you must add this user name to the domain administrator group.

BPCIN0221E The user name *userID* is not a part of the local administrator group.

Explanation

The user name must be in the local administrator group to install IBM Spectrum Control.

Action

Add the user name to the local administrator group and continue with the installation.

BPCIN0222E The current login user *loginuser* is not an administrator or a member of the local administrator group. The installation program cannot continue.

Explanation

This user name is not an administrator or a member of the local administrator group.

Action

To continue the installation program, you must add this user to the local administrator group.

BPCIN0223E The current login user *loginuser* is not a member of the local Db2 administrator group. The installation program cannot start.

Explanation

This user name is not an administrator or a member of the domain administrator group.

Action

To start the installation program, you must add this user name to the local Db2 administrator.

BPCIN0224E The user *userID* is a domain account and cannot be used to log in to Db2. You must enter a separate user name for Db2.

Explanation

This error message is displayed if you specify a domain user account for the common user when you install IBM Spectrum Control, but a local user account was previously created during the Db2 installation. The error occurs when you click Next on the Single Server Installation Information page during the installation of IBM Spectrum Control.

Action

To resolve this error and gain access to Db2, complete the following steps:

1. On the Single Server Installation Information page, click Configure Database repository.
2. On the Configure the Database Repository page, enter the Db2 user name without a Windows domain name prefix, enter the password for the Db2 user, and click Validate.

3. When the validation is complete, click OK to return to the Single Server Installation Information page.
4. Click Next on the Single Server Installation Information page to continue the installation.

BPCIN0225E The current logged in user *loginuser* does not have Db2 SYSADM authority. To provide the user with Db2 SYSADM authority, log in by using a user name with SYSADM authority and run the following commands: `db2cmd db2set -g DB2_GRP_LOOKUP=local,TOKENLOCAL db2 force application all db2stop db2start`

Explanation

The current user is not in an operating system group with Db2 SYSADM authority.

Action

The current user is not in an operating system group with Db2 SYSADM authority. For more information on acquiring user group information in the Windows operating system, go to the IBM Knowledge Center at <http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome> and search for "access token" For more information about the `sysadm_group` parameter, go to the IBM Knowledge Center at <http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome> and search for "sysadm_group"

BPCIN0226E Login as a windows domain user in order to install IBM Spectrum Control using a windows domain account.

Explanation

Login as a windows domain user in order to install IBM Spectrum Control using a windows domain account.

Action

Login as a windows domain user in order to install IBM Spectrum Control using a windows domain account.

BPCIN0229E The Monitoring Agent for Windows OS - Watchdog and Monitoring Agent for Windows OS - Primary services must be stopped before you can continue. After you upgrade, you must restart these services.

Explanation

The Monitoring Agent for Windows OS - Watchdog and Monitoring Agent for Windows OS - Primary services are running on Windows and can cause the upgrade to fail.

Action

To stop the Monitoring Agent for Windows OS - Watchdog and Monitoring Agent for Windows OS - Primary services, click Start > Control Panel > Administrative Tools > Services.

BPCIN0230E The Monitoring Agent for Windows OS - Watchdog and Monitoring Agent for Windows OS - Primary services must be stopped before you can continue. After you uninstall IBM Spectrum Control, you must restart these services if the required reboot is postponed.

Explanation

The Monitoring Agent for Windows OS - Watchdog and Monitoring Agent for Windows OS - Primary services are running and can cause the IBM Spectrum Control uninstallation to fail.

Action

To stop the Monitoring Agent for Windows - Watchdog and Monitoring Agent for Windows - Primary services, click Start > Control Panel > Administrative Tools > Services.

BPCIN0231E Tivoli Common Reporting is not installed on your system. You must install it to continue with the IBM Spectrum Control installation.

Explanation

Jazz for Service Management is already installed. In order to continue with the IBM Spectrum Control installation and to have Cognos BI Reports installed, you must install Tivoli Common Reporting.

Action

Install Tivoli Common Reporting and reinstall IBM Spectrum Control.

BPCIN0233E Jazz for Service Management is not installed in the specified location. Reenter the correct installation location for Jazz for Service Management.

Explanation

Reenter the correct installation location for Jazz for Service Management.

Action

Reenter the correct installation location for Jazz for Service Management and reinstall IBM Spectrum Control.

BPCIN0234E The Jazz for Service Management user credentials that you entered were incorrect. Check the user credentials and try again.

Explanation

The Jazz for Service Management user credentials that you entered were incorrect.

Action

Check the user credentials and try again.

BPCIN0235W The Jazz for Service Management and Tivoli Common Reporting servers are not running. To start the servers, run the following command: *file* It takes a few minutes for these servers to start up and initialize. Wait for a few minutes before resuming the installation.

Explanation

Jazz for Service Management is not running.

Action

Start the Jazz for Service Management server.

BPCIN0236W The Tivoli Common Reporting server is not running. If you have already started the server , wait for a few more minutes. It takes a while for the server to start up and initialize. Otherwise, you can start the server by running this command: *file*

Explanation

Start the the Tivoli Common Reporting server.

Action

The Tivoli Common Reporting server is not running.

BPCIN0238E The Tivoli Common Reporting server at *tcrlocation* cannot be reached. If the Jazz for Service Management server has not completed startup, wait for a few more minutes and then click OK. It takes a while for the server to start and initialize. Otherwise, if the Jazz for Service Management server is started but Tivoli Common Reporting still cannot be reached, restart the Jazz for Service Management server using the following commands: *stopServer -username username -password password startServer*

Explanation

An error occurred connecting to the Tivoli Common Reporting server.

Action

Restart Tivoli Common Reporting by restarting the Jazz for Service Management server. To verify Tivoli Common Reporting is reachable, run the following command: *trcmd -checkstatus -reportingEngine [-minutestowait minutes]*. You must authenticate with *-user user -password password*.

BPCIN0239E The Tivoli Common Reporting server configuration cannot be exported.

Explanation

An error occurred while exporting the Tivoli Common Reporting server configuration.

Action

Restart the Tivoli Common Reporting server.

BPCIN0240E The Tivoli Common Reporting server configuration file "*configuration file*" was not created.

Explanation

An error occurred while exporting the Tivoli Common Reporting server configuration.

Action

Restart the Tivoli Common Reporting server.

BPCIN0241E The Tivoli Common Reporting configuration cannot be upgraded.

Explanation

An error occurred while upgrading the Tivoli Common Reporting server configuration.

Action

Restart the Tivoli Common Reporting server.

BPCIN0242E The Windows registry check indicates that .NET 3.5 is not available on this Windows 2012 server. You must install .NET 3.5 to continue. To install .NET 3.5 type the following commands in a 64 bit windows powershell: `Import-Module ServerManager Add-WindowsFeature as-net-framework`

Explanation

You must install .NET 3.5 to continue.

Action

To install .NET 3.5, type the following commands in a 64 bit windows powershell: `Import-Module ServerManager Add-WindowsFeature as-net-framework`

BPCIN0244E An error occurred while enumerating the local administrator group membership. On the current computer, on the Properties page of this group, remove the user names that are displayed with a security ID (SID). An example of a SID is S-1-5-21-337177553-1671989427-887411491-500 and is used instead of a user name.

Explanation

An error occurred while enumerating the local administrator group membership. On the current computer, on the Properties page of this group, remove the user names that are displayed with a SID. An example of a SID is S-1-5-21-337177553-1671989427-887411491-500. This string is used instead of a user name and is a reference or a link to a domain user who no longer exists on the domain controller computer.

Action

An error occurred while enumerating the local administrator group membership. On the current computer, on the Properties page of this group, you must remove the user names that are displayed with a SID. An example of a SID is S-1-5-21-337177553-1671989427-887411491-500. This string is used instead of a user name and is a reference or a link to a domain user who no longer exists on the domain controller computer.

BPCIN0245E An error occurred while enumerating the local Db2 administrator group membership. On the current computer, on the Properties page of this group, remove the user names that are displayed with a security ID (SID). An example of a SID is S-1-5-21-337177553-1671989427-887411491-500 and is used instead of a user name.

Explanation

An error occurred while enumerating the local Db2 administrator group membership. On the current computer, on the Properties page of this group, remove the user names that are displayed with a security ID (SID). An example of a SID is S-1-5-21-337177553-1671989427-887411491-500. This string is used instead of a user name

and is a reference or link to a domain user name that no longer exists on the domain controller computer.

Action

An error occurred while enumerating the local Db2 administrator group membership. On the current computer, on the Properties page of this group, remove the user names that are displayed with a security ID (SID). An example of a SID is S-1-5-21-337177553-1671989427-887411491-500. This string is used instead of a user name and is a reference or link to a domain user name that no longer exist on the domain controller computer.

BPCIN0246E Tivoli Storage Productivity Center *version is installed. Before you can upgrade to version 5, you must upgrade Tivoli Storage Productivity Center to version 4.*

Explanation

You cannot upgrade Tivoli Storage Productivity Center version 4.1 directly to version 5.2. You must first upgrade version 4.1 to version 4.2x or 5.1, then upgrade version 4.2x or 5.1 to version 5.2.

Action

For information about upgrading to version 4.2x, see the Tivoli Storage Productivity Center Knowledge center version 4.2.2. Search for Upgrading and migrating the Tivoli Storage Productivity Center family. After you upgrade to version 4.2x, upgrade to version 5.x.

BPCIN0247E An error occurred during the domain check prevalidation process. Verify that the domain controller computer is available and then restart this domain member machine.

Explanation

An error occurred during the domain check prevalidation process. Verify that the domain controller computer is available and then restart this domain member machine.

Action

An error occurred during the domain check prevalidation process. Verify that the domain controller computer is available and then restart this domain member machine.

BPCIN0248E The IBM Spectrum Control installation program supports only fully qualified user names on Windows domain member machines. Specify the user name *userID* by using the "Domain_name\\" or the "Machine_name\\" prefix. The detected domain name is "*domainName*".

Explanation

The installation program supports only fully qualified user names on Windows domain member machines. Specify the user name by using the "Domain_name\\" or the "Machine_name\\" prefix.

Action

Enter a valid user name and continue with the installation.

BPCIN0250W Before you can upgrade the database repository, you must first stop the Data server, Device server, Alert server, Export server, and Web GUI server on the remote server. For more information about stopping IBM Spectrum Control servers use the following link.

Explanation

You must stop the servers before upgrading the database repository.

Action

Scripts to stop and start servers are provided in SC_installation_directory/scripts.

BPCIN0252E Db2 has been installed by using a domain user account. The IBM Spectrum Control installation software does not allow the Db2 user name *userID* because this user name exists in both the Windows domain and the local operating system repositories. Specify a Db2 user name that exists only in the Windows Domain registry.

Explanation

The installation software does not allow a Db2 user name that exists in both the Windows domain and local operating system repositories.

Action

Specify a Db2 user name that exists only in the Windows Domain registry.

BPCIN0253W All IBM Spectrum Control components are installed.
GUIModeMessage

Explanation

Action

BPCIN0254W IBM Spectrum Control servers are installed on *remoteHost*.

Explanation

The IBM Spectrum Control servers are already installed on the remote computer.

Action

Click quit to exit the installation program.

BPCIN0255E The database *dbName* must be upgraded to the current version. For more information about upgrading Db2, go to the IBM Knowledge Center and search for "Upgrading Db2"

Explanation

The Db2 database must be upgraded to the current version that is supported.

Action

BPCIN0256W If you continue upgrading to IBM Spectrum Control, the reports you have from Tivoli Storage Productivity Center Version *oldVersion* will be deleted. When you uninstall Tivoli Integrated Portal, the Authentication Services Server is also uninstalled. If you have storage subsystems that are configured to use LDAP

authentication through the Authentication Services Server, before you upgrade the product and uninstall Tivoli Integrated Portal, reconfigure the storage subsystems so that these subsystems do not use LDAP authentication. Click OK to continue or Cancel to select another option.

Explanation

Action

BPCIN0257E The Db2 database installation on *dbPath* path was made when the creation of 8.3 filenames was disabled on this server. The IBM Spectrum Control installation will fail when Db2 is installed to a path with spaces and 8.3 filenames were not enabled when it was installed.

Explanation

The installation may fail.

Action

Install Db2 in a folder that does not have a blank space in the name and run the IBM Spectrum Control installation program again.

BPCIN0258E The LDAP configuration export did not succeed. For more information, review the log files, and in the LDAP export command output, search for "exportLDAPRepositories".

Explanation

The LDAP configuration export did not succeed because of exceptions when running the WebSphere wsadmin command.

Action

The LDAP configuration export did not succeed. For more information, review the log files. In the LDAP export command output, search for "exportLDAPRepositories outputStream cmd result" or for "exportLDAPRepositories errorStream cmd result".

BPCIN0260E The connection to the local database failed because the user name or password is not valid.

Explanation

You must specify a valid user name and password to connect to the local database.

Action

Check your local database for a valid user name and password.

BPCIN0261E The installation program was unable to connect to the local database. Check the server to verify that you can connect to Db2. For more information about Db2 connection issues, go to the IBM Knowledge Center at <http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome>.

Explanation

An error occurred when the installation program tried to create a JDBC connection to the local database. A reason for this error might be that the database was not started on the server.

Action

Check the server and verify that you can connect to Db2. For more information about Db2 connection issues, go to the IBM Knowledge Center at <http://www-01.ibm.com/support/knowledgecenter/SSEPGG/welcome> and search for "Troubleshooting Db2 Server."

BPCIN0262E The IBM Spectrum Control upgrade process has stopped because Jazz for Service Management and Tivoli Integrated Portal are using the same ports. To continue, you must install Jazz for Service Management and Tivoli Common Reporting again and use ports that are different from the ports that are used by Tivoli Integrated Portal.

Explanation

The upgrade has been stopped because Jazz for Service Management and Tivoli Integrated Portal are using the same ports.

Action

The IBM Spectrum Control upgrade process has stopped because Jazz for Service Management and Tivoli Integrated Portal are using the same ports. To continue, you must install Jazz for Service Management and Tivoli Common Reporting again and use ports that are different from the ports that are used by Tivoli Integrated Portal."

BPCIN0264E The version of the database repository component cannot be queried from the Db2 Database Manager. SQL Error Message is *sqlerrmsg*

Explanation

The version of the database repository cannot be queried from Db2 Database Manager.

Action

Check your Db2 instance configuration.

BPCIN0265E An error occurred during the prevalidation of the component because the *file_name* is missing.

Explanation

The IBM Spectrum Control installation stopped during prevalidation because the file is missing. This file is required to complete the prevalidation.

Action

To resolve this error, locate the file and run the installation program again. If the problem persists, contact IBM Software Support.

BPCIN0266E Db2 has been installed by using a domain user name. The user name *dbuser* is not a member of the domain Db2 administrator group *domainDB2AdminGroup*, so the IBM Spectrum Control installation program cannot start.

Explanation

Db2 has been installed by using a domain user name, but this user name is not a domain administrator or a member of the domain Db2 administrator group.

Action

To start the IBM Spectrum Control installation program, you must add this user name to the domain Db2 administrator group.

BPCIN0267E The password *password* cannot start with the following special characters: *characters*.

Explanation

The password cannot be used because it starts with special characters that are invalid.

Action

Enter a valid password.

BPCIN0268E The installation program could not stop the .You must manually stop the by running the script.

Explanation

The installation program could not stop the server.

Action

Try to stop the server manually and continue with the installation or upgrade.

BPCIN0269E The Jazz(tm) for Service Management installation image Version "*retrivedVersion*" is not at the correct build level.

Explanation

The Jazz(tm) for Service Management installation image is not at the correct build level. Download and install the latest build.

Action

The Jazz(tm) for Service Management installation image is not at the correct build level. For more information about accessing the latest Jazz(tm) for Service Management build, follow the instructions in the IBM Spectrum Control download document.

BPCIN0272E "*installDirectory*" cannot be found where the Jazz(tm) for Service Management installation files are extracted. Extract the Jazz(tm) for Service Management installation files to a local directory before you start the Jazz(tm) for Service Management installation program.

Explanation

A directory or file is missing from the local directory where you extracted the Jazz(tm) for Service Management installation files.

Action

After you extract the Jazz(tm) for Service Management installation files to a local directory, enter this local directory path on the "Install Jazz(tm) for Service Management" page.

If you cannot resolve the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0275E A 32-bit version of Db2 has been detected. IBM Spectrum Control can only be installed with a 64-bit version of Db2. Install a supported 64-bit version of Db2, and install IBM Spectrum Control again.

Explanation

IBM Spectrum Control can only be installed with a 64-bit version of Db2.

Action

Install a supported 64-bit version of Db2 and install IBM Spectrum Control again.

BPCIN0276E A required system library could not be loaded. Review the log files for more information about this error.

Explanation

All required system libraries must be installed on AIX or Linux systems.

Action

If you are running on AIX or Linux, go to the IBM Knowledge Center and search on "Starting Installation". Follow the instructions on installing the required libraries then run the installation program again.

BPCIN0277W The Jazz(tm) for Service Management installation image that you selected is not the latest version. Click Yes to continue with the current version and not install or upgrade reports. Click No to install or upgrade to Jazz(tm) for Service Management Version "*latest_JazzSM_available_version*" and Tivoli Common Reporting Version "*latest_TCR_available_version*".

Explanation

The Jazz(tm) for Service Management installation image that you selected is not the latest version. Download and install or upgrade to the latest version.

Action

The Jazz(tm) for Service Management installation image that you selected is not the latest version. For more information about installing or upgrading to the latest version of Jazz(tm) for Service Management and Tivoli Common Reporting, go to the IBM Knowledge Center and search for "Install Jazz".

BPCIN0279W IBM Spectrum Control does not support the version of Jazz(tm) for Service Management that is installed on this computer. You must upgrade to Jazz(tm) for Service Management *minJazzSMVersion* before you can continue.

Explanation

Before you can install or upgrade Cognos BI Reports, a supported version of Jazz(tm) for Service Management must be installed on the system.

Action

For information about the Jazz(tm) for Service Management supported version, go to the IBM Spectrum Control download document. Install or upgrade the required Jazz(tm) for Service Management version and run the installation program again.

BPCIN0281E The current version of the Tivoli Common Reporting installation image is not at the correct level.

Explanation

The current version of the Tivoli Common Reporting installation image is not at the correct level. Download and install the latest version.

Action

The current version of the Tivoli Common Reporting installation image is not at the correct level. For information about installing the latest version of Jazz(tm) for Service Management, see the instructions in the IBM Spectrum Control download document.

BPCIN0282E The installation package for Tivoli Common Reporting Version 3.1.0.1 cannot be found on your computer. To continue installing IBM Spectrum Control, download the installation package for Tivoli Common Reporting Version 3.1.0.1 in the same directory where you downloaded the installation package for Tivoli Common Reporting Version 3.1.0.2.

Explanation

To generate Cognos BI Reports in IBM Spectrum Control, you must install Jazz for Service Management Version 1.1.0.3. To install Jazz for Service Management Version 1.1.0.3, you must download the installation packages for Tivoli Common Reporting Version 3.1.0.2 and Version 3.1.0.1. Tivoli Common Reporting 3.1.0.2 is a Fix Pack and cannot be installed by itself.

Action

To continue installing IBM Spectrum Control, download the installation package for Tivoli Common Reporting Version 3.1.0.1 in the same directory where you downloaded the installation package for Tivoli Common Reporting Version 3.1.0.2.

BPCIN0283E You cannot upgrade "*latest_version*" to the lower version "*lower_version*".

Explanation

You must download a higher version of the product.

Action

To continue upgrading, download a higher version of the product.

BPCIN0284E The Jazz for Service Management installation directory does not have execution rights. You can add execution rights to the directory by running the following command:
`chmod -R u+x "JazzSM_build_folder"`

Explanation

The Jazz for Service Management installation directory does not have execution rights. You can add execution rights to the directory by running the following command:
`chmod -R u+x JazzSM_build_folder`

Action

After you extract the Jazz(tm) for Service Management installation files to a local directory, check that the directory and the files have execution rights

If you cannot resolve the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0285E The `getDB2Inst.sh "user_name"` command is displaying a "null" result because the Db2 profile was not sourced for user "user_name" or the environment for "user_name" has been corrupted.

Explanation

The `getDB2Inst.sh` command is displaying a "null" result for one of the following reasons:

- You did not source the Db2 profile before the installation of IBM Spectrum Control.
- The environment for the specified user has been corrupted.

To resolve this issue, run the `"su - user -c "echo $DB2INSTANCE" "` command and ensure that the correct Db2 instance name is displayed.

Action

The `getDB2Inst.sh` command is displaying a "null" result for one of the following reasons:

- You did not source the Db2 profile before the installation.
- The environment for the specified user has been corrupted.

To resolve this issue, run the `"su - user -c "echo $DB2INSTANCE" "` command and ensure that the correct Db2 instance name is displayed.

BPCIN0286W The operation of the storage resource agent is limited on the operating system of the local server. Total Disk Space and Available Disk Space on this server cannot be determined.

Explanation

The storage resource agent will not scan all local volumes.

Action

Only a subset of local storage data will be available.

BPCIN0287E Directory *directory* is located on a memory based file system (RAM disk) and cannot be used for installing IBM Spectrum Control.

Explanation

The directory that is selected for the installation is located on a memory based file system (RAM disk). Everything in this folder is deleted after a system reboot.

Action

Select a directory that is not located on a memory based file system (RAM disk).

BPCIN0288E The upgrade process detected an error with the previously installed version of the product. The installation directory of the previously installed version of the product contains corrupted files. Remove or fix the corrupted files and run the upgrade process again. See the `lax*-out.txt` and `lax*-err.txt` log files in the system temporary directory for details of the error.

Explanation

The previously installed version of the product has some corrupted configuration files. See the lax*-out.txt and lax*-err.txt log files, in %temp% on Windows operating systems or in /tmp on AIX or Linux operating systems, for details. Also, check the TPC.log file in the installation directory of the previous version of the product for more information. Remove or fix the corrupted files to fix the upgrade error and run the upgrade process again.

Action

If you cannot fix the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0289E You must install Db2 before you can install Cognos BI Reports. Ensure that Db2 is already installed on the system and the db2profile is sourced.

Explanation

Db2 must be installed on the local system and you must source the db2profile if you want to install Cognos BI Reports.

Action

No further information is available

BPCIN0290E The IBM Spectrum Control installation image is corrupted. Extract the IBM Spectrum Control installation image again into an empty directory.

Explanation

This message is for informational purposes only.

Action

This message is for informational purposes only.

BPCIN0291E You need to define 'localhost' in the 'hosts' file that is used by Jazz for Service Management and restart the host system.

Explanation

Jazz for Service Management cannot start if localhost is not defined correctly in the 'hosts' file.

Action

This message is for informational purposes only.

BPCIN0293E The IBM Spectrum Control installation program cannot find the license folder. Ensure the license folder is extracted into "extractorDir".

Explanation

The installation cannot continue until the license folder is available in the IBM Spectrum Control installation folder.

Action

Extract the license folder from the downloaded installation files into the IBM Spectrum Control installation folder.

BPCIN0295E Tivoli Storage Productivity Center *old_version* is installed. Before you upgrade to IBM Spectrum Control version *new_version*, you must upgrade to Tivoli Storage Productivity Center version 5.2.7.

Explanation

You cannot upgrade directly from Tivoli Storage Productivity Center 4.2.X or earlier to IBM Spectrum Control 5.2.8 or later. Upgrade to version 4.2.2, then upgrade to Tivoli Storage Productivity Center version 5.2.7.

Action

For information about upgrading to version 4.2.2, go to the product documentation at http://www-01.ibm.com/support/knowledgecenter/SSNE44_4.2.2/com.ibm.tpc_V422.doc/fqz0_t_upgrading_all.html. After you upgrade to version 4.2.2, upgrade to 5.2.7.

BPCIN0296E The installation program could not rename the *short_name* directory because the folder is in use by other processes. Quit the installation program and stop all processes that access the directory *full_path*.

Explanation

An error occurred while renaming a directory.

Action

Stop all processes that are accessing the directory and start the installation program again.

BPCIN0297E The *user_name* user does not have the full control permission for the *folder_path* folder.

Explanation

The user who is running the installation program must have the full control permission for the specified folder.

Action

Assign the full control permission to the folder for the user who is logged on to run the installation program.

BPCIN0298E The *user_name* fenced user does not have full permissions for the *directory_path* database directory.

Explanation

The fenced user must have read, write, and execute permissions for the directory.

Action

Use the `chmod` command to assign full permissions to the directory for the fenced user.

BPCIN0299E The installation location "*install_location*" has a naming conflict with a file or folder named *folder_path*. Rename

the *folder_path* file or folder.

Explanation

The installation location causes conflicts on the Windows operating system with an existing file or folder.

Action

Rename the specified file or folder.

BPCIN0300E The value *db2_variable_existing_value* is not allowed for the Db2 variable DB2_LIMIT_FENCED_GROUP. Use the `db2set -g DB2_LIMIT_FENCED_GROUP=OFF` command to change the value of the variable to OFF.

Explanation

The value ON or YES is not allowed for the variable DB2_LIMIT_FENCED_GROUP.

Action

Use the `db2set -g DB2_LIMIT_FENCED_GROUP=OFF` command to change the value of the variable DB2_LIMIT_FENCED_GROUP to OFF.

BPCIN0301E The *versionFile* file is empty. Extract IBM Spectrum Control again and make sure the *version.txt* file contains a valid build string.

Explanation

The build failed to add a version string to the *version.txt* file.

Action

Extract IBM Spectrum Control again and make sure the *version.txt* file contains a valid build string.

BPCIN0302E Db2 Advanced Enterprise Server Edition is not supported. Install Db2 Enterprise Server Edition and then proceed with your installation.

Explanation

Before you can install IBM Spectrum Control, a supported version of Db2 Enterprise Server Edition must be installed.

Action

Install a supported version of Db2 Enterprise Server Edition and run the installation program again. For a list of Db2 Enterprise Server Edition versions that are supported, go to <http://www-01.ibm.com/support/docview.wss?uid=swg27039833#Databases>.

BPCIN0303W Your storage environment has switches with obsolete data sources. Configure up-to-date data sources after you upgrade.

Explanation

The upgrade process identified switches with one of the following obsolete data sources:

- Brocade switches that are managed by using SNMP.
- Cisco switches that are managed by using a Storage Resource agent or an SMI agent.

Action

Configure up-to-date data sources after you upgrade IBM Spectrum Control. To monitor switches, provide the following credentials:

- For Brocade switches that use a version of Fabric OS earlier than 8.2.1, provide credentials for a Brocade Network Advisor.
- For Brocade switches that use Fabric OS 8.2.1 or later, provide credentials for the switches.
- For Cisco switches, provide SNMP credentials.

Related reference

- [Find the Supported Hardware, Products and Platforms Interoperability Matrix Links](#)

BPCIN0304E You cannot upgrade from a Basic Edition license.

Explanation

The Basic Edition license is no longer supported and upgrades with the Basic Edition license are not supported. Upgrade your current product license and then upgrade to IBM Spectrum Control. Please contact IBM Software Support for more details.

Action

The Basic Edition license is no longer supported and upgrades with the Basic Edition license are not supported. Upgrade your current product license and then upgrade to IBM Spectrum Control.

BPCIN0306W The Storage Resource agent registered successfully with the Data server but some problems occurred after the registration. Review the log files in the *SRA_log_name* directory for additional information.

Explanation

The Storage Resource agent registered successfully with the Data server but there were some problems after the registration.

Action

Review the Storage Resource agent log files for more information. If the problem persists, contact IBM Software Support.

BPCIN0307W You are currently connected to devices using a CIM interface. This upgrade might require new certificates to be generated on your CIM managed devices to continue monitoring them. To resolve this issue, go to <http://www.ibm.com/support/docview.wss?uid=swg21976237>

Explanation

To create new certificates for your CIM managed device go to <http://www.ibm.com/support/docview.wss?uid=swg21976237>

Action

To create new certificates for your CIM managed device go to <http://www.ibm.com/support/docview.wss?uid=swg21976237>

BPCIN0307E The security certificates for the Web server have expired. Renew the security certificates and run the installation program again. See the *lax*-out.txt* and *lax*-err.txt* log files in the system temporary directory for details of the error and how to resolve the issue.

Explanation

The trust.p12 and key.p12 keystores in the directories and the WebServerNode for the WAS WebServer might contain the expired certificates.

Action

Renew the security certificates and run the installation program again. See the lax*-out.txt and lax*-err.txt log files in the system temporary directory for details of the error and how to resolve the issue. You can search for areInvalidWASSecurityCertificates in the latest lax*-out.txt to see details about the error and how to renew the certificates.

BPCIN0308W Brocade switches that use SNMP services were detected. SNMP data sources are no longer used to manage Brocade fabrics and switches. Instead, you must use the embedded SMI agent in Brocade Network Advisor.

Explanation

You must use the embedded SMI agent in Brocade Network Advisor to collect metadata about switches after you upgrade to IBM Spectrum Control V5.2.10 or later. Brocade switches that used SNMP services before the upgrade will display as unreachable. To collect metadata about those switches, you must configure them to use the SMI agent in Brocade Network Advisor. The metadata that was collected for these switches before you upgraded to V5.2.10 is retained.

Important: The embedded SMI agent is only available in the Professional Plus and Enterprise editions of Brocade Network Advisor. If you don't already have a license for Brocade Network Advisor, you can purchase it from Broadcom before February 2020 and use it as an SMI agent for IBM Spectrum Control. After that time, it will no longer be available for purchase, although it will be supported by Broadcom until 2022. For more information about end of support for Brocade Network Advisor, see <https://www.broadcom.com/support/fibre-channel-networking/eol>.

Action

Follow these steps to configure your Brocade switches to use the Brocade Network Advisor SMI agent:

- If you do not have Brocade Network Advisor, download and install the Brocade version. Select the SMI Agent Only option during installation.
- In the IBM Spectrum Control GUI, go to Network -> Switches.
- Select one or more switches and select Actions-> Connection-> Modify Connection.
- Enter the switch information and save.
- Go to Actions -> Data Collection > Schedule to schedule data collections for the switches.

BPCIN0309E The port configuration export did not succeed. For more information, review the log files, and in the port export command output, search for "exportCurrentPorts".

Explanation

The port configuration export did not succeed because of exceptions when running the WebSphere wsadmin command.

Action

The port configuration export did not succeed. For more information, review the log files. In the port export command output, search for "exportCurrentPorts outputStream cmd result" or for "exportCurrentPorts errorStream cmd result".

BPCIN0310E Spectrum Control installation requires that the Administrators group in Windows is assigned "Debug programs" privilege. Check this setting, log out of and back into Windows, and try again.

Explanation

"Debug programs" rights in the Windows security policy are not assigned to the Administrators group. Check the current value. "Debug programs" can be found under Administrative Tools > Local Security Policy > Local Policies > User Rights Assignment.

Action

Assign the "Debug programs" right to the Administrators group, log out and log back in again, and run the installation program again. See the lax*-out.txt and lax*-err.txt log files in the system temporary directory for further details of the error and how to resolve the issue.

BPCIN0311E Spectrum Control installation requires the wmic command to work. Check this command, ensure the service "Windows Management Instrumentation" is running and try again.

Explanation

The wmic command does not work because the service "Windows Management Instrumentation" is not running.

Action

Check if the wmic command works ok and if the service "Windows Management Instrumentation" is running. You can run "wmic os get osarchitecture" from a command line as Administrator to see the error message. See the lax*-out.txt and lax*-err.txt log files in the system temporary directory for details of the error and how to resolve the issue.

BPCIN0312E Spectrum Control installation requires that wmic command to work. Check this command, ensure service "Windows Management Instrumentation" is working, disable any antivirus and try again.

Explanation

wmic command does not work because service "Windows Management Instrumentation" is not working or an antivirus is running.

Action

Check if wmic command and service "Windows Management Instrumentation" are working and disable any antivirus. You can run as Administrator "wmic os get osarchitecture" from a command line to see the error message. See the lax*-out.txt and lax*-err.txt log files in the system temporary directory for details of the error and how to resolve the issue.

BPCIN0313E Spectrum Control installation requires that wmic command to work. Check this command and the PATH environment variable and try again.

Explanation

wmic command does not work because either does not exist or the folder that contains it is not in the PATH environment variable.

Action

Check if wmic command works. You can run as Administrator "wmic os get osarchitecture" from a command line to see the error message. See the lax*-out.txt and lax*-err.txt log files in the system temporary directory for details of the error and how to resolve the issue.

BPCIN0314E Spectrum Control installation requires that wmic command to work. Check this command and try again.

Explanation

wmic command does not work because the exit code is not successful.

Action

Check if wmic command works. You can run as Administrator "wmic os get osarchitecture" from a command line to see the error message and then "echo %ERRORLEVEL%" to see the exit code. See the lax*-out.txt and lax*-err.txt log files in the system temporary directory for details of the error and how to resolve the issue.

BPCIN0315E Spectrum Control installation requires that chcp command to work. Check this command and try again.

Explanation

chcp command does not work because the exit code is not successful.

Action

Check if chcp command works. You can run as Administrator "chcp" from a command line to see the error message and then "echo %ERRORLEVEL%" to see the exit code. See the lax*-out.txt and lax*-err.txt log files in the system temporary directory for details of the error and how to resolve the issue.

BPCIN0316E The IBM Spectrum Control upgrade process does not support custom materialized query tables (MQTs). Remove any MQTs from the database and run the upgrade process again.

Explanation

No further information is available.

Action

See the lax*-out.txt and lax*-err.txt log files in the system temporary directory for details of the error and how to resolve the issue.

BPCIN0317E The IBM Spectrum Control upgrade process does not support missing vendor information from the .com.zerog.registry.xml file. Add the missing information and run the upgrade process again. See the lax*-out.txt and lax*-err.txt log files to resolve the issue.

Explanation

No further information is available.

Action

See the lax*-out.txt and lax*-err.txt log files in your system temporary directory for details on the error and how to resolve the issue.

BPCIN0318E The IBM Spectrum Control upgrade process does not support duplicated users user in IBM WebSphere Application Server related to the Web server. Remove the duplicate user from a federated repository other than the localOS default repository; run the upgrade process again.

Explanation

No further information is available.

Action

One possible cause is that the LDAP repository is defined in the WebSphere Application Server WebServer that contains the duplicate user. See the lax*-out.txt and lax*-err.txt log files in the temporary directory of your system for details on the error and how to resolve the issue.

BPCIN0319E The currently installed version of AIX XL C/C++ RUNTIME is not supported. You must upgrade to a supported level. The

version that was detected is *version*.

Explanation

IBM Spectrum Control cannot be installed on this system because the version of AIX XL C/C++ RUNTIME is not supported. You must upgrade AIX XL C/C++ RUNTIME to a supported level to continue.

Action

For information about version of AIX XL C/C++ RUNTIME, go to <http://www-01.ibm.com/support/docview.wss?uid=swg21975779>.

BPCIN0320E PAM (Pluggable Authentication Modules) isn't installed on your system. Installing PAM would resolve the issue.

Explanation

No further information is available.

Action

See the `lax*-out.txt` and `lax*-err.txt` log files in your system temporary directory for details on the error and how to resolve the issue.

BPCIN0321E An error occurred when restoring authorities before creating the database *dbName*. Review the log files for more information.

Explanation

The installation program cannot recover from the installation error.

Action

Review the log files in the following directory for an explanation of the error:

- For Windows operating systems: `TPC_installation_directory\logs\`
- For UNIX or Linux operating system: `TPC_installation_directory/logs/`

`TPC_installation_directory` is where IBM Spectrum Control is installed.

The log files are:

- `msgTPCInstall.log`
- `traceTPCInstall.log`

Correct the problem. If you cannot correct the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0323E The Data Server port has different values in the configuration files. Verify that the port value is correct in the `installDir/data/config/server.config` file, and the `installDir/config/InstallVariable.properties` file. Also, verify that the `PORT_NUMBER` column where `SERVER_TYPE` is `serverType` is correct in the `T_RES_Server` database table.

Explanation

Verify that the correct port number value is the same in the files and the data base table.

Action

Add the correct port number value in the files and the data base table and run the installer program again. If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0324W The current certificate used to secure the connection of the IBM Spectrum Control Data server and the Storage Resource agents cannot be replaced with the latest certificate that provides higher security. After you complete the upgrade, you can create new certificates by following the information in the following link.

Explanation

Action

Update the Data server keystore to use SHA256 signing algorithm by following the steps documented in following link.
"https://www.ibm.com/support/knowledgecenter/SS5R93_5.3.6/com.ibm.spectrum.sc.doc/fqz0_r_create_custom_certificate_ssl.html"

Related reference

- [Replacing default SSL certs with custom certs](#)

BPCIN0325E IBM Tivoli Storage Productivity Center *old_version* is installed. Before you can upgrade to IBM Spectrum Control *new_version*, you must upgrade to IBM Spectrum Control 5.3.0 or later. For more information about upgrading to IBM Spectrum Control 5.4.0 or later, go to the following link:

Explanation

You cannot directly upgrade from IBM Tivoli Storage Productivity Center 5.2.17.4 or earlier to IBM Spectrum Control 5.4.0 or later. You must upgrade to 5.3.0 or later, then upgrade to 5.4.0 or later.

Action

For information about upgrading to IBM Spectrum Control 5.4.0, go to the product documentation at https://www.ibm.com/support/knowledgecenter/SS5R93_5.4.0/ and search for "Preparing for an upgrade."

BPCIN0328E The IBM Spectrum Control installation requires the `LD_LIBRARY_PATH` environment variable be set. Set this environment variable and try the installation again. If you ran `setup.bin` from the IBM Spectrum Control installation directory with `sudo` and the `LD_LIBRARY_PATH` variable was properly set, then try the `sudo -E LD_LIBRARY_PATH=$LD_LIBRARY_PATH ./setup.bin` command.

Explanation

The IBM Spectrum Control installation does not work because an environment variable is not properly set.

Action

Verify that the environment variables are properly set. See the `lax*-out.txt` and `lax*-err.txt` log files in the system temporary directory for details of the error and how to resolve the issue.

BPCIN0329E The IBM Spectrum Control installation requires that the `LIBPATH` environment variable be set. Set this environment variable and try the installation again.

Explanation

The IBM Spectrum Control installation does not work because an environment variable is not properly set.

Action

Verify that the environment variables are properly set. See the `lax*-out.txt` and `lax*-err.txt` log files in the system temporary directory for details of the error and how to resolve the issue.

BPCIN0330E An SQL exception was created when querying the Db2 Database Manager. The SQL error message is: *sqlerrmsg*.

Explanation

The required data cannot be queried from Db2 Database Manager.

Action

Verify that your Db2 instance is configured correctly. If you cannot resolve the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCIN0331E The *tpcregFile* file is not valid because it contains a failed or partial upgrade to a previous IBM Spectrum Control Version *newTPCVersion*. Continue with the incomplete upgrade before you try to upgrade to the new IBM Spectrum Control Version *varNewUpgradeVersion*.

Explanation

Continue with the incomplete upgrade before you try to upgrade to the new IBM Spectrum Control Version.

Action

Check the `.tpcreg` file for more details and review the installation log files.

Review the `.tpcreg` file in one of the following directories:

- For Windows operating systems: `TPC_installation_directory_uninst\`
- For AIX or Linux operating systems: `TPC_installation_directory/_uninst/`

`TPC_installation_directory` is where IBM Spectrum Control is installed.

See the `lax*-out.txt` and `lax*-err.txt` log files, in `%temp%` on Windows operating systems or `/tmp` on AIX or Linux operating systems, for details.

BPCIN0335E The IBM Spectrum Control installation requires you install the `libstdc++` package on the AIX operating system. Download and install the `libstdc++` package and start the installation again. You can download the package here:

<https://www.ibm.com/developerworks/aix/library/aix-toolbox/alpha.html>.

Explanation

Action

BPCIN0336E The version of the `libstdc++` package you installed is lower than the minimum required Version *minVersion*. Download and upgrade the `libstdc++` package and start the installation again. You can download the package here:
<https://www.ibm.com/developerworks/aix/library/aix-toolbox/alpha.html>.

Explanation

Action

BPCIN0337E The *componentServer* password failed to validate. Check that you entered the password correctly and try again.

Explanation

Action

BPCIN0338E The *componentServers* passwords failed to validate. Check that you entered the passwords correctly and try again.

Explanation

Action

BPCIN0339W Your current Db2 instance is using a trial license. The license will expire on *expireDate*. Once your trial expires Db2 will not start. If you want to upgrade your Db2 license now, use the information in the following link.

Explanation

Action

Related reference

- [🔗 Licensing Db2](#)

BPCIN0340W Your current Db2 instance is using a trial license. The license will expire on *expireDate*. Once your trial expires Db2

will not start. If you want to upgrade your Db2 license now, goto [link](#).

Explanation

Action

Related reference

- [🔗 Licensing Db2](#)

BPCIN0343W There are new certificate requirements that are strictly enforced for macOS Catalina users which might affect their ability to access the IBM Spectrum Control GUI. During an upgrade of IBM Spectrum Control, certificates self-signed by IBM Spectrum Control will be made compliant automatically. However, if one or more of your certificates are not self-signed by IBM Spectrum Control, see the following links for more information. Validate that your certificates are compliant.

Explanation

Action

Related reference

- [🔗 IBM Spectrum Control and macOS Catalina \(10.15\) Increased Security Policies on SSL Certificates](#)

BPCIN0344E Your current Db2 instance trial license has expired. If you want to upgrade your Db2 license now, goto the following link.

Explanation

Action

Related reference

- [🔗 Licensing Db2](#)

BPCIN0345E The user user is part of the "Deny access to this computer from the network" security policy setting on your computer. Either contact your administrator to have the user removed from the security policy or enter a different user. For more information on required user privileges in installation scenarios, go to the following link.

Explanation

Action

Related reference

- [🔗 Required user privileges in installation scenarios](#)

BPCIN0346E Your current Db2 License type: Community is not supported for IBM Spectrum Control. To upgrade your Db2 license, go to the following link:

Explanation

Action

Related reference

- [🔗 Licensing Db2](#)

BPCRE - Alert server messages

- [BPCRE0001I Trying again to connect to the repository database...](#)
- [BPCRE0003E An invalid type was specified for a day range.](#)
- [BPCRE0007E Error adding rule rule_name.](#)
- [BPCRE0008I Drools Version: version_number.](#)
- [BPCRE0009I Duration - build rules:value seconds.](#)
- [BPCRE0010I Duration - setup KIE session: value seconds.](#)
- [BPCRE0011E Failure writing rule rule to database.](#)
- [BPCRE0012E Failure inserting data for rule rule to database.](#)
- [BPCRE0013E Failure rebuilding rules after membership change in group.](#)

BPCRE0001I Trying again to connect to the repository database...

Explanation

The Alert server is trying again to connect to the repository database...

BPCRE0003E An invalid type was specified for a day range.

Explanation

Invalid type for day range. Valid day range types are either days of a week or dates of a month.

Action

Specify a valid day range and try again.

BPCRE0007E Error adding rule *rule_name*.

Explanation

An internal error has occurred when constructing the rule.

Action

Contact IBM Software Support.

Related reference

- [🔗 Getting support](#)

BPCRE0008I Drools Version: *version_number*.

Explanation

The current Drools Version.

BPCRE0009I Duration - build rules: *value* seconds.

Explanation

Duration - build rules.

BPCRE0010I Duration - setup KIE session: *value* seconds.

Explanation

Duration - setup KIE session.

BPCRE0011E Failure writing rule *rule* to database.

Explanation

Rules engine encountered an error writing a rule to the database.

Action

Contact IBM Software Support.

Related reference

- [Getting support](#)
-

BPCRE0012E Failure inserting data for rule *rule* to database.

Explanation

The Rules engine encountered an error inserting data to a rule.

Action

Contact IBM Software Support.

Related reference

- [Getting support](#)
-

BPCRE0013E Failure rebuilding rules after membership change in *group*.

Explanation

The Rules engine encountered an error rebuilding rules affected by a group or application membership change.

Action

Contact IBM Software Support.

Related reference

- [Getting support](#)

BPCRS - Spectrum Control Middleware messages

- [BPCRS0000E A required key *key* is not found in map.](#)
- [BPCRS0001E Value *value* is not found in list.](#)
- [BPCRS0002E Some or all of the required key-value bindings are not found. Required Keys: *required*. Supplied keys: *supplied*.](#)
- [BPCRS0003E key value *value* is not in list.](#)
- [BPCRS0004E Expected non-null value for *var*.](#)
- [BPCRS0005E Expected non-empty string value for *str*.](#)
- [BPCRS0006E Value supplied for *name* is not of type *type*.](#)
- [BPCRS0007I The log file retention settings were successfully updated.](#)
- [BPCRS0008E The log file retention settings were not successfully updated.](#)
- [BPCRS0009I The alert disposition settings were updated.](#)
- [BPCRS0010E The alert disposition settings were not updated.](#)
- [BPCRS0011I The history and log retention settings were successfully updated.](#)
- [BPCRS0012E The history and log retention settings were not successfully updated.](#)
- [BPCRS0013I Step *currentStep* of totalSteps : *stepName*](#)
- [BPCRS0014I *stepName* completed](#)
- [BPCRS0015E *stepName* failed](#)

BPCRS0000E A required key *key* is not found in map.

Explanation

A required key is not found in the request object.

Action

BPCRS0001E Value *value* is not found in list.

Explanation

Value is not in the list of accepted values.

Action

BPCRS0002E Some or all of the required key-value bindings are not found. Required Keys: *required*. Supplied keys: *supplied*.

Explanation

Some or all of the required keys are not present.

Action

BPCRS0003E key value *value* is not in list.

Explanation

Key value is not in the list of accepted values.

Action

BPCRS0004E Expected non-null value for *var*.

Explanation

Expected non-null value.

Action

BPCRS0005E Expected non-empty string value for *str*.

Explanation

Expected non-empty string.

Action

BPCRS0006E Value supplied for *name* is not of type *type*.

Explanation

Expected a value that is of a specific type.

Action

BPCRS0007I The log file retention settings were successfully updated.

Explanation

The log file retention settings were successfully updated.

Action

BPCRS0008E The log file retention settings were not successfully updated.

Explanation

The log file retention settings were not successfully updated.

Action

BPCRS0009I The alert disposition settings were updated.

Explanation

The alert disposition settings were updated.

BPCRS0010E The alert disposition settings were not updated.

Explanation

The alert disposition settings were not updated.

Action

Check the values that you just entered in the Alert Disposition panel. Look for any invalid settings, correct those settings, and try to save the alert disposition settings again.

BPCRS0011I The history and log retention settings were successfully updated.

Explanation

The history and log retention settings were successfully updated.

Action

BPCRS0012E The history and log retention settings were not successfully updated.

Explanation

The history and log retention settings were not successfully updated.

Action

BPCRS0013I *Step currentStep of totalSteps : stepName*

Explanation

Action

BPCRS0014I *stepName* completed

Explanation

Action

BPCRS0015E *stepName* failed

Explanation

Action

BPCCM - Data collector messages

- [BPCCM0001I Data collection is being performed by hostname.](#)

BPCCM0001I Data collection is being performed by *hostname*.

Explanation

The data collector installed on the machine at the specified hostname will collect data for the subsystem. The host must have connectivity to the device.

BPCDP - Data processor messages

- [BPCDP0000I Performance data for natural key resource at date and time timestamp was collected and processed successfully.](#)
- [BPCDP0001E Error while collecting and processing performance data for natural key resource at date and time timestamp. Performance data was not collected and processed.](#)
- [BPCDP0002E The processing of performance data for the resource could not be completed.](#)
- [BPCDP0003E No performance data is available at the current time for this resource.](#)
- [BPCDP0004I Performance data was retrieved and persisted but aggregation of data to higher-level components didn't complete because the relationship to the higher-level components couldn't be determined.](#)
- [BPCDP0005E Could not save the performance data that was collected from the resource.](#)
- [BPCDP0006I Performance data at date and time timestamp was processed and saved successfully for the resource.](#)
- [BPCDP0007E The resource is missing. The resource is required.](#)
- [BPCDP0008E Identifying information for the resource is missing. This information is required.](#)
- [BPCDP0009E Information identifying the resource type is invalid: system type .](#)
- [BPCDP0010E Information uniquely identifying the resource is missing. This information is required.](#)
- [BPCDP0011E Information uniquely identifying the resource is invalid.](#)
- [BPCDP0012E The UUID for the tenant's resource is invalid.](#)
- [BPCDP0013E The start time for the performance data is invalid: start time](#)
- [BPCDP0014E The end time for the performance data is invalid: end time](#)
- [BPCDP0015I Performance data at date and time timestamp was processed and saved successfully for the resource, but the data processing raised warnings.](#)
- [BPCDP0016W Performance data for natural key resource at date and time timestamp was collected and processed successfully, but the data processing raised warnings.](#)

BPCDP0000I Performance data for *natural key* resource at *date and time* timestamp was collected and processed successfully.

Explanation

The Data Processor successfully completed performance monitoring for the resource. The performance data was collected at the indicated resource time stamp in the server time zone. The saved information was either received from the resource or was computed based on the information that was received from the resource.

Action

No action is required.

BPCDP0001E Error while collecting and processing performance data for *natural key* resource at *date and time* timestamp. Performance data was not collected and processed.

Explanation

The Data Processor failed to complete performance monitoring for the resource. The performance data was not collected for the indicated resource time stamp in the server time. The information was not received from the resource or computed by the Data Processor based on the information that was received from the resource.

Action

No action is required.

BPCDP0002E The processing of performance data for the resource could not be completed.

Explanation

The current attempt to process performance data did not complete. No performance data is persisted for this time period. The collection of performance data that is running is not affected by this error.

Related reference

- [!\[\]\(a43b62a38b6e2844e794f4301a08d3ba_img.jpg\) Getting support](#)

BPCDP0003E No performance data is available at the current time for this resource.

Explanation

The data collector did not provide any performance data at the current time.

Related reference

- [Getting support](#)

BPCDP0004I Performance data was retrieved and persisted but aggregation of data to higher-level components didn't complete because the relationship to the higher-level components couldn't be determined.

Explanation

Performance statistics for some high-level components are aggregated from base component statistics. For example, pool statistics are aggregated from volumes belonging to the pool, and the volume to pool relationship is needed to do such an aggregation. The relationship data might not be available if the device probe hasn't yet completed. When this relationship can't be determined, statistics for the pool can't be aggregated.

BPCDP0005E Could not save the performance data that was collected from the resource.

Explanation

The data processor was unable to persist the performance data for this time period.

BPCDP0006I Performance data at *date and time* timestamp was processed and saved successfully for the resource.

Explanation

The performance data was collected at the indicated time. The saved information was either received from the data collector, or was computed based on the information received from the data collector.

Action

None. No action is required.

BPCDP0007E The resource is missing. The resource is required.

Explanation

The data collector did not provide a required property to the data processor. Without this property, the data processor can't process performance data for this time period.

Related reference

- [Getting support](#)

BPCDP0008E Identifying information for the resource is missing. This information is required.

Explanation

The data collector did not provide a required property to the data processor. Without this property, the data processor can't process performance data for this time period.

Related reference

- [Getting support](#)

BPCDP0009E Information identifying the resource type is invalid: *system type* .

Explanation

The data collector provided an invalid property to the data processor. Without this property, the data processor can't process performance data for this time period.

Related reference

- [Getting support](#)

BPCDP0010E Information uniquely identifying the resource is missing. This information is required.

Explanation

The data collector did not provide a required property to the data processor. Without this property, the data processor can't process performance data for this time period.

Related reference

- [Getting support](#)

BPCDP0011E Information uniquely identifying the resource is invalid.

Explanation

The data collector provided an invalid property to the data processor. Without this property, the data processor can't process performance data for this time period.

Related reference

- [Getting support](#)

BPCDP0012E The UUID for the tenant's resource is invalid.

Explanation

The data collector provided an invalid property to the data processor. Without this property, the data processor can't process performance data for this time period.

Related reference

- [Getting support](#)

BPCDP0013E The start time for the performance data is invalid: *start time*

Explanation

The data collector provided an invalid property to the data processor. Without this property, the data processor can't process performance data for this time period.

Related reference

- [Getting support](#)

BPCDP0014E The end time for the performance data is invalid: *end time*

Explanation

The data collector provided an invalid property to the data processor. Without this property, the data processor can't process performance data for this time period.

Related reference

- [Getting support](#)

BPCDP0015I Performance data at *date and time* timestamp was processed and saved successfully for the resource, but the data processing raised warnings.

Explanation

The Data Processor successfully completed performance monitoring for the resource. The performance data was collected at the indicated resource time stamp in the server time zone. The saved information was either received from the resource or was computed based on the information that was received from the resource. The warnings indicate that either some performance information wasn't received from the resource, or couldn't be computed from the information that was received.

Action

None. No action is required.

BPCDP0016W Performance data for *natural key* resource at *date and time* timestamp was collected and processed successfully, but the data processing raised warnings.

Explanation

The Data Processor successfully completed performance monitoring for the resource. The performance data was collected at the indicated resource time stamp in the server time zone. The saved information was either received from the resource or was computed based on the information that was received from the resource. The warnings indicate that either some performance information wasn't received from the resource, or couldn't be computed from the information that was received.

Action

No action is required.

BPCSS - Scheduler messages

- [BPCSS0000E An error occurred while collecting performance data from the device. The collection is being attempted by a different collector.](#)
- [BPCSS0001W The data collection is taking longer than expected.](#)
- [BPCSS0002E Currently, there is no data collector available for this device.](#)
- [BPCSS0005I Performance monitor is starting at an interval of interval interval units. This action was requested by user name.](#)
- [BPCSS0008I Collection can no longer continue due to invalid credentials. Use the 'Modify Connection' dialog to fix the storage system credentials and resume collection.](#)
- [BPCSS0009E Failed to save the performance monitor schedule.](#)
- [BPCSS0010E A job cannot be run for resource resourceName because there is a job already running for the resource.](#)
- [BPCSS0011W The schedule change was saved but the update to the active collection did not happen.](#)
- [BPCSS0012I Performance monitor is stopped. This action was requested by user name.](#)
- [BPCSS0013I Performance monitor is stopped.](#)
- [BPCSS0014I Performance monitor is starting at an interval of interval interval units.](#)
- [BPCSS0015I Performance monitor collection interval was updated to interval interval units. This action was requested by user name.](#)
- [BPCSS0016I Performance monitor collection interval was updated to interval interval units.](#)

- [BPCSS0017I Performance monitor is enabled. This action was requested by user name.](#)
- [BPCSS0018I Performance monitor is enabled.](#)
- [BPCSS0019I Performance monitor is disabled. This action was requested by user name.](#)
- [BPCSS0020I Performance monitor is disabled.](#)
- [BPCSS0021W Performance monitor is starting. The initial attempt to start collection failed so it is retried. This action was requested by user name.](#)
- [BPCSS0022W Performance monitor is starting. The initial attempt to start collection failed so it is retried.](#)
- [BPCSS0023I Performance monitor collection interval is enabled and updated to interval interval units. This action was requested by user name.](#)
- [BPCSS0024I Performance monitor collection interval is enabled and updated to interval interval units.](#)
- [BPCSS0025E Access to the agent or device is denied. Ensure that valid credentials are specified for agent agent name.](#)
- [BPCSS0026E New performance data is not yet available for the device. Statistics with time stamps later than time stamp could not be found.](#)
- [BPCSS0027E The performance monitor failed due to an internal error.](#)
- [BPCSS0028E The value that is specified as parameter \(value\) is invalid.](#)
- [BPCSS0029E Cannot connect to the device with the address IP address.](#)
- [BPCSS0030E Cannot connect to the SNMP data source IP address.](#)
- [BPCSS0031E Cannot authenticate with the provided user credentials.](#)
- [BPCSS0032E Passphrase is incorrect for subsystem param1.](#)
- [BPCSS0033E Passphrase is required. Specify one for subsystem param1.](#)
- [BPCSS0034E Verify that they private key that was provided for subsystem param1 was in the OpenSSH file format. If it is in another format, it needs to be converted before it can be used.](#)
- [BPCSS0035E The user does not have the required authority to complete the task or command.](#)
- [BPCSS0036E Cannot connect to the storage system or cluster.](#)
- [BPCSS0037W The device cannot be reached.](#)
- [BPCSS0038E The device or device agent did not respond within the allotted time.](#)
- [BPCSS0039E The host name or IP address {Q} is not valid.](#)
- [BPCSS0040E The host name or IP address is not valid.](#)
- [BPCSS0041E Cannot connect to the device.](#)
- [BPCSS0042E Cannot connect to the SNMP data source.](#)
- [BPCSS0043E Passphrase is incorrect.](#)
- [BPCSS0044E Passphrase is required.](#)
- [BPCSS0045E Access to the device is denied. Ensure that valid credentials are specified.](#)
- [BPCSS0046E Verify that they private key that was provided was in the OpenSSH file format. If it is in another format, it needs to be converted before it can be used.](#)
- [BPCSS0047E New performance data is not yet available for the device.](#)
- [BPCSS0048E The parameter for the Performance Manager API is invalid.](#)
- [BPCSS0049E Schedule is not enabled for the resource resource.](#)
- [BPCSS0050W Performance data could not be collected for device device name because the device or data source cannot be reached \(reason reason code\). The current samples are skipped.](#)
- [BPCSS0051E The device or device agent did not respond within the allotted time \(timeout valuesseconds\).](#)
- [BPCSS0052W Performance data continuity is broken. The device might have been reset or rebooted. record count performance data records were discarded.](#)
- [BPCSS0053W No valid performance data was provided by the monitored resource. Zero performance data records were inserted into the database.](#)
- [BPCSS0054E A timeout occurred while polling the performance statistics for this device: device name](#)
- [BPCSS0055E Performance data was not collected for device device name due to error error trace. The current samples are skipped.](#)
- [BPCSS0056E The last performance Data Collection was not readable for device device name, the collection failed with error error trace.](#)
- [BPCSS0057E Cannot connect to the switch with the provided IP address, host name, protocol, and port.](#)
- [BPCSS0058E Cannot authenticate to the switch with the provided user name and password.](#)
- [BPCSS0059E The specified user name does not have the required permissions for the switch.](#)
- [BPCSS0112I The probe was stopped. This action was requested by user name.](#)
- [BPCSS0113I The probe is stopped.](#)
- [BPCSS0114I The probe is starting at an interval of interval interval units.](#)
- [BPCSS0115I The probe interval was updated to interval. This change was requested by user name.](#)
- [BPCSS0116I The probe interval was updated to interval.](#)
- [BPCSS0117I The probe is enabled. This action was requested by user name.](#)
- [BPCSS0118I The probe is enabled.](#)
- [BPCSS0119I The probe is disabled. This action was requested by user name.](#)
- [BPCSS0120I The probe is disabled.](#)
- [BPCSS0123I The Probe collection interval is enabled and updated to interval. This action was requested by user name.](#)
- [BPCSS0124I Probe interval is enabled and updated to interval.](#)
- [BPCSS0105I Probe is starting at an interval of interval interval units. This action was requested by user name.](#)

BPCSS0000E An error occurred while collecting performance data from the device. The collection is being attempted by a different collector.

Explanation

There are multiple reasons why this condition could be encountered. Generally the error is something that could be specific to one collector and not occur on other collectors.

Some examples are as follows:

1. The data collector cannot resolve the hostname of the subsystem.
2. Network communication between the data collector and the device is blocked.

Action

If this problem persists and it is not caused by the data collector being down, consider reducing the data collection frequency.

BPCSS0001W The data collection is taking longer than expected.

Explanation

The amount of time that has elapsed since the last performance sample was collected and processed is longer than expected considering the selected collection interval.

This error might be a result of all collectors that can contact the device are down. Or because the amount of time required to collect and process the data is more than the collection interval.

Action

Confirm that the data collectors are running and are able to connect to the server. Try reducing the collection frequency. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BPCSS0002E Currently, there is no data collector available for this device.

Explanation

No data collector has been found that has successfully been able to connect to the device.

Action

Verify that the device is running and fully operational.

If it is in a hung state, it might be necessary to reboot the device.

Ensure that the IP address and hostname are correct and that the user has the required permissions to complete the task.

BPCSS0005I Performance monitor is starting at an interval of *interval interval units*. This action was requested by *user name*.

Explanation

Performance statistics are collected from the device.

BPCSS0008I Collection can no longer continue due to invalid credentials. Use the 'Modify Connection' dialog to fix the storage system credentials and resume collection.

Explanation

Access is denied by the storage system with the credentials currently specified.

The credentials usually consist of a username and password, but can also encompass other security related parameters such as ssh keys, depending on the type of device being accessed and, where applicable, the access method selected.

Action

Go to the Block Storage Systems overlay, and right click on the storage system.

Hover over 'Connections' and select 'Modify Connection' to fix the credentials for the device.

Ensure that the user entered has the correct permissions for collection, and that the password, passphrase, or ssh key file for the device are correct.

BPCSS0009E Failed to save the performance monitor schedule.

Explanation

For more information about the cause of the error, see the trace logs.

Action

Try updating the performance monitor schedule again. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BPCSS0010E A job cannot be run for resource *resourceName* because there is a job already running for the resource.

Explanation

A new job cannot be run for a schedule definition if a previous job from that schedule definition did not finish running. Only one running job per schedule at a time is permitted.

Action

Stop the job or wait for the job to finish and try again.

BPCSS0011W The schedule change was saved but the update to the active collection did not happen.

Explanation

The schedule changes take effect the next time the job is started.

Action

Restart the job to force the update.

BPCSS0012I Performance monitor is stopped. This action was requested by *user name*.

Explanation

The performance monitor is stopped and no performance data are collected for this device until the performance monitor is restarted.

Action

Restart the performance monitor, as needed.

BPCSS0013I Performance monitor is stopped.

Explanation

The performance monitor is stopped and no performance data are collected for this device until the performance monitor is restarted.

Action

Restart the performance monitor, as needed.

BPCSS0014I Performance monitor is starting at an interval of *interval interval units*.

Explanation

Performance statistics are collected from the device.

BPCSS0015I Performance monitor collection interval was updated to *interval interval units*. This action was requested by *user name*.

Explanation

Performance statistics are collected at the new interval.

BPCSS0016I Performance monitor collection interval was updated to *interval interval units*.

Explanation

Performance statistics are collected at the new interval.

BPCSS0017I Performance monitor is enabled. This action was requested by *user name*.

Explanation

Performance statistics are collected.

BPCSS0018I Performance monitor is enabled.

Explanation

Performance statistics are collected.

BPCSS0019I Performance monitor is disabled. This action was requested by *user name*.

Explanation

Performance statistics collection is disabled.

BPCSS0020I Performance monitor is disabled.

Explanation

Performance statistics collection is disabled.

BPCSS0021W Performance monitor is starting. The initial attempt to start collection failed so it is retried. This action was requested by *user name*.

Explanation

Performance monitor started but the collector failed to start the collection. The server continues to look for a collector that can perform the collection.

BPCSS0022W Performance monitor is starting. The initial attempt to start collection failed so it is retried.

Explanation

Performance monitor started but the collector failed to start the collection. The server continues to find a collector that can perform the collection.

BPCSS0023I Performance monitor collection interval is enabled and updated to *interval interval units*. This action was requested by *user name*.

Explanation

Performance statistics are collected at the new interval.

BPCSS0024I Performance monitor collection interval is enabled and updated to *interval interval units*.

Explanation

Performance statistics are collected at the new interval.

BPCSS0025E Access to the agent or device is denied. Ensure that valid credentials are specified for agent *agent name*.

Explanation

When retrieving performance statistics for the device, access is denied by the device or the device agent.

Therefore, no performance statistics could be retrieved.

No performance data is inserted for the device in this time period.

The next performance data sample that is recorded into the database might represent an average over more than the configured interval length.

Action

Ensure that the correct agent address is specified. Ensure that the specified credentials are valid and allow access to the device or agent for performance data collection.

The credentials usually consist of a user name and password, but can also encompass other security-related parameters such as ssh keys or authentication tokens.

The type of credentials used depends on the type of device or agent that is being accessed and, where applicable, the access method selected.

BPCSS0026E New performance data is not yet available for the device. Statistics with time stamps later than *time stamp* could not be found.

Explanation

In case performance data is cached by the device or device agent, the performance manager ensures that the most recently retrieved performance statistics are indeed newer than the previously retrieved statistics.

If the time stamp in the message is "null", then no statistics were previously retrieved, and the performance manager is unable to get ANY statistics for the device.

No performance data is inserted for the device in this time period. The next performance data sample that is recorded into the database might represent an average over more than the configured interval length.

Action

Ensure that the device and device agent are fully operational.

It might be necessary to restart either device or agent, if it is in a hung state.

Ensure that if the device has multiple clocks (for example for multiple nodes or controllers), that the clocks are synchronized to within a few minutes.

BPCSS0027E The performance monitor failed due to an internal error.

Explanation

If the error is recoverable, the performance manager attempts to restart.

This failure might be temporary, and can result in a temporary disruption of performance data collection for the device.

Action

From the Resources menu, click Storage Systems.

Check the status of the performance monitors.

If the performance monitor is stopped, right-click the storage system, click Data Collection, and then click Start Performance Monitor.

If you still can't complete the action, contact IBM Support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BPCSS0028E The value that is specified as parameter (value) is invalid.

Explanation

The parameter value passed as the argument on the Performance Manager API call is invalid. The method called returns this error indicator.

Action

Refer to the method or class documentation to determine the allowed values for the particular parameter.

Modify the caller of this method to pass a valid value to the API.

Related reference

- [IBM Spectrum Control REST API](#)

BPCSS0029E Cannot connect to the device with the address *IP address*.

Explanation

The storage device with this IP address cannot be reached.

Action

Ensure that the IP address of the storage device is correct and that the device is running properly.

If the problem persists, contact IBM Support

Related reference

- [📄 Default locations of log files](#)

BPCSS0030E Cannot connect to the SNMP data source *IP address*.

Explanation

A test connection cannot be established to the data source at the specified IP address.

This error might occur if the data source is not available or the SNMP community is entered incorrectly. The SNMP community name is shared by one or more SNMP hosts and is used to authenticate messages that are received by those hosts.

BPCSS0031E Cannot authenticate with the provided user credentials.

Explanation

The user name or password entered for the device is not correct.

Action

Make sure that the user name and password are correct for the device that is being added.

Reenter the user name or password and click Add again.

If you are adding a Storwize V7000 Unified storage system, specify the IP address of the block component, not the filer component.

BPCSS0032E Passphrase is incorrect for subsystem *param1*.

Explanation

The passphrase for the truststore is incorrect.

Action

Provide the correct passphrase.

BPCSS0033E Passphrase is required. Specify one for subsystem *param1*.

Explanation

The passphrase for the truststore is missing.

Action

Provide the missing passphrase.

BPCSS0034E Verify that the private key that was provided for subsystem *param1* was in the OpenSSH file format. If it is in another format, it needs to be converted before it can be used.

Explanation

Action

BPCSS0035E The user does not have the required authority to complete the task or command.

Explanation

The level of authority that is required for the task or command depends on the type of resource that the user is managing.

This error can occur in the following situations:

- For IBM SONAS resources, if the user is an rssh restricted account and can issue only a restricted set of commands on the resource.
- For IBM Spectrum Scale, if the user that is used to log on to the cluster node does not have privileges to monitor the GPFS cluster. The user must have root privileges on the cluster node or have privileges to run a set of specified administration commands using the sudo command. For information about monitoring IBM Spectrum Scale without requiring root privileges, go to the IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/search/tpch_t_configuring_sudo_access?scope=SSQRB8.

Action

Ensure that the user has the required permissions to complete the task.

Related reference

- [🔗 Required user roles for monitoring resources](#)

BPCSS0036E Cannot connect to the storage system or cluster.

Explanation

The device that is being added might not be up and running.

Also, the IP address, host name, or user credentials entered for the storage system or cluster might not be valid.

Action

Verify that the device is up and running.

Also, make sure that the IP address, host name, and user credentials are correct for the device that is being added. Reenter the values and click Add again.

BPCSS0037W The device cannot be reached.

Explanation

The current attempt to retrieve a set of performance data from the device failed.

No performance data is inserted for the device in this time period.

The next performance data sample that is recorded into the database might represent an average over more than the configured interval length.

The immediate operation of the running performance monitor is unaffected.

Action

If the problem persists for an hour or longer, ensure that your device and data source (if applicable) are operational.

Also, ensure that a network path exists between the IBM Spectrum Control server and the device or data source, including any appropriate firewall pass-throughs.

If nothing is wrong with the device or device agent, or with the network path, try canceling and restarting the performance monitor job.

If the problem is still not resolved, contact IBM Support.

Related reference

- [🔗 Getting support](#)

BPCSS0038E The device or device agent did not respond within the allotted time.

Explanation

When retrieving performance statistics for the device, the requested performance data was not returned before the timeout expired.

Either the device or agent is unresponsive, or is much slower than expected.

No performance data is inserted for the device in this time period.

The next performance data sample that is recorded into the database might represent an average over more than the configured interval length.

Action

Ensure that the device and device agent are fully operational.

It might be necessary to restart either device or agent, if it is in a hung state.

If you have reason to believe that the device and agent are operational but are slower than expected, you can also attempt to increase the timeout value that is used by the performance manager, which is set in configuration file `device/conf/pm.conf`.

BPCSS0039E The host name or IP address {0} is not valid.

Explanation

The IP address or host name entered for the device is not valid.

Action

Make sure that the IP address and host name are valid for the device being added.

Reenter the IP address or host name and click Add again.

BPCSS0040E The host name or IP address is not valid.

Explanation

The IP address or host name entered for the device is not valid.

Action

Make sure that the IP address and host name are valid for the device being added.

Reenter the IP address or host name and click Add again.

BPCSS0041E Cannot connect to the device.

Explanation

The storage device with this IP address cannot be reached.

Action

Ensure that the IP address of the storage device is correct and that the device is running properly.

If the problem persists, contact IBM Support

Related reference

- [Default locations of log files](#)

BPCSS0042E Cannot connect to the SNMP data source.

Explanation

A test connection cannot be established to the data source at the specified IP address.

This error might occur if the data source is not available or the SNMP community is entered incorrectly.

The SNMP community name is shared by one or more SNMP hosts and is used to authenticate messages that are received by those hosts.

BPCSS0043E Passphrase is incorrect.

Explanation

The passphrase for the truststore is incorrect.

Action

Provide the correct passphrase.

BPCSS0044E Passphrase is required.

Explanation

The passphrase for the truststore is missing.

Action

Provide the missing passphrase.

BPCSS0045E Access to the device is denied. Ensure that valid credentials are specified.

Explanation

When retrieving performance statistics for the device, access is denied by the device or the device agent.

Therefore, no performance statistics could be retrieved. No performance data is inserted for the device in this time period.

The next performance data sample that is recorded into the database might represent an average over more than the configured interval length.

Action

Ensure that the correct agent address is specified, and that the specified credentials are valid and allow access to the device or agent for performance data collection.

The credentials usually consist of a user name and password, but can also encompass other security-related parameters such as ssh keys or authentication tokens, depending on the type of device or agent that is being accessed and, where applicable, the access method selected.

BPCSS0046E Verify that the private key that was provided was in the OpenSSH file format. If it is in another format, it needs to be converted before it can be used.

Explanation

Action

BPCSS0047E New performance data is not yet available for the device.

Explanation

In case performance data is cached by the device or device agent., the performance manager ensures that the most recently retrieved performance statistics are indeed newer than the previously retrieved statistics.

If the time stamp in the message is "null", then no statistics were previously retrieved, and the performance manager is unable to get ANY statistics for the device.

No performance data is inserted for the device in this time period.

The next performance data sample that is recorded into the database might represent an average over more than the configured interval length.

Action

Ensure that the device and device agent are fully operational.

It might be necessary to restart either device or agent, if it is in a hung state.

Ensure that if the device has multiple clocks (for example for multiple nodes or controllers), that the clocks are synchronized to within a few minutes.

BPCSS0048E The parameter for the Performance Manager API is invalid.

Explanation

The parameter value passed as argument on the Performance Manager API call is invalid.

The method called returns this error indicator.

Action

Refer to the method or class documentation to determine the allowed values for the particular parameter.

Modify the caller of this method to pass a valid value to the API.

BPCSS0049E Schedule is not enabled for the resource *resource*.

Explanation

Schedule is not enabled. Cannot be started.

Action

In order to be started the job should be enabled first.

BPCSS0050W Performance data could not be collected for device *device name* because the device or data source cannot be reached (reason *reason code*). The current samples are skipped.

Explanation

The current attempt to retrieve a set of performance data from the device failed.

No performance data is inserted for the device in this time period.

The next performance data sample that is recorded into the database might represent an average over more than the configured interval length.

The immediate operation of the running performance monitor is unaffected.

The reason code can be used to help identify the exact cause of the problem encountered:

- Reason Code 0 indicates that the exact reason for the failure could not be determined.
- Reason Code 1 indicates a bad target (device or data source) address. This condition can occur when the user-specified host name or IP address, or the target port number are invalid such that they would cause the formation of an invalid URL or IP address. This type of failure is rare and can usually also be identified by a `java.net.UnknownHostException` printed in the trace logs.
- Reason Code 2 indicates the problem to be an unknown target address. This condition can occur when a host name rather than an IP address was specified as target address, and:
 - either the network is down,
 - the specified host name cannot be resolved (for example, the name server cannot be contacted, or the name server is down, or the specified host name is not known to the name server), or
 - the specified host name can be resolved by the name server but no longer exists on the network.

This type of failure can usually also be identified by a `java.net.UnknownHostException` printed in the trace logs.

- Reason Code 3 indicates the problem to be an unreachable target address. This condition can occur when an IP address rather than a host name is specified as target address, and either the network or a part of the network is down or is blocked by a firewall (the host cannot be contacted), or the specified IP address does not exist on the network. This type of failure can usually also be identified by a `java.net.NoRouteToHostException` in the trace logs.
- Reason Code 4 indicates the problem to be an unresponsive target. This condition can occur when the target server is powered off, or when the server is not listening on the port, which is the target of the communication. For example, when the web server or SMI-S provider is not operational. This type of failure can

usually also be identified by a `java.net.ConnectException` printed in the trace logs.

- Reason Code 5 indicates a communication timeout for communication that uses UDP rather than TCP, for example when you use SNMP data sources. This condition can occur when the target server cannot be reached, or when the SNMP data source is disabled on the target server, or when the SNMP port (161) is blocked by a firewall.

Action

If the problem persists for an hour or longer, ensure that your device and data source (if applicable) are operational.

Also, ensure that a network path exists between the IBM Spectrum Control server and the device or data source, including any appropriate firewall pass-throughs.

If nothing is wrong with the device or device agent, or with the network path, try canceling and restarting the performance monitor job.

If the problem is still not resolved, contact IBM Support.

Related reference

- [Getting support](#)

BPCSS0051E The device or device agent did not respond within the allotted time (*timeout value*seconds).

Explanation

When retrieving performance statistics for the device, the requested performance data did not return before the timeout expired.

Either the device or agent is unresponsive, or is much slower than expected.

No performance data is inserted for the device in this time period.

The next performance data sample that is recorded into the database might represent an average over more than the configured interval length.

Action

Ensure that the device and device agent are fully operational.

It might be necessary to restart either device or agent, if it is in a hung state.

If you have reason to believe that the device and agent are operational but are slower than expected, you can also attempt to increase the timeout value that is used by the performance manager, which is set in configuration file `device/conf/pm.conf`.

BPCSS0052W Performance data continuity is broken. The device might have been reset or rebooted. *record count* performance data records were discarded.

Explanation

This message indicates that invalid performance information was received from the device.

In general, performance information is represented as a set of ever-increasing counters, and actual statistics are computed by taking the difference between two consecutive sets of such counters.

However if a counter appears to have decreased instead of increased between consecutive sets, the information is unusable and is discarded.

Note that counters can be expected to decrease if they are reset to zeroes, which might happen normally when a device is reset or rebooted (for example when new firmware is loaded), or in some cases when a device agent (CIMOM for example) is reset or rebooted.

In those situations, this warning message can be safely ignored.

If this warning occurs when no reset or reboot has occurred, the device or device agent might be generating incorrect performance statistics, and you might have to contact your device vendor for further instructions.

Action

Determine if the device or device agent has been reset or rebooted.

Those situations include loading of new firmware, or fail-over and fail-back scenarios for ESS, DS6000, and DS8000 storage subsystems.

In any of these cases, the reset of performance counters is expected behavior, and this warning message can be safely ignored.

If not one of these cases, the device might be generating incorrect performance data, which might or might not lead to inaccurate performance reports.

Contact your device vendor for further instructions in those cases.

BPCSS0053W No valid performance data was provided by the monitored resource. Zero performance data records were inserted into the database.

Explanation

The performance monitor contacted the resource and tried to collect data. However, the resource did not provide valid performance counter information.

If message HWNPM2124W is also displayed, the resource was able to provide performance data, but it was determined to be invalid and discarded.

Action

The operation of the performance monitor is not affected, and it will attempt to retrieve performance data again for the next sample interval.

If the monitored resource continues to provide no performance data, ensure it is fully operational.

Where appropriate, also ensure that performance functionality is enabled for the resource.

BPCSS0054E A timeout occurred while polling the performance statistics for this device: *device name*

Explanation

The performance collection did not complete in specified time.

Action

The operation of the performance monitor is not affected, and it will attempt to retrieve performance data again for the next sample interval. If the monitored resource continues to receive timeout, ensure that it is fully operational. If appropriate, also ensure that performance functionality is enabled for the resource.

BPCSS0055E Performance data was not collected for device *device name* due to error *error trace*. The current samples are skipped.

Explanation

The performance collection encountered an unknown error.

Action

The operation of the performance monitor is not affected, and it will attempt to retrieve performance data again for the next sample interval. If the monitored resource continues to receive timeout, ensure that it is fully operational. If appropriate, also ensure that performance functionality is enabled for the resource.

BPCSS0056E The last performance Data Collection was not readable for device *device name*, the collection failed with error *error trace*.

Explanation

The read operation of the last collected performance data failed.

Action

The last collected performance data might be corrupted. Restart the Data Collection or add the device again.

BPCSS0057E Cannot connect to the switch with the provided IP address, host name, protocol, and port.

Explanation

One or more of the following details for the switch are incorrect: IP address, host name, protocol, port.

Action

Enter the correct IP address, host name, protocol, and port for the switch. Make sure that the switch meets the minimum version requirements and is configured to use the specified protocol.

BPCSS0058E Cannot authenticate to the switch with the provided user name and password.

Explanation

One or more of the following details for the switch are incorrect: User name, password.

Action

Enter the correct user name and password for the switch.

BPCSS0059E The specified user name does not have the required permissions for the switch.

Explanation

The user name is a valid user on the switch, but does not have the chassis-role permission.

Action

Specify another user name that has the chassis-role permission. Alternatively, enter another user name and password for the switch.

BPCSS0112I The probe was stopped. This action was requested by user *name*.

Explanation

Metadata will not be collected about the device until the probe is restarted.

Action

Restart the probe, as needed. To start a probe, complete these steps:

1. From the menu bar, go to the resource list page for a resource type. For example, to start a probe for a block storage system, select Resources > Block Storage Systems.
2. Right-click the resource and select Data Collection > Start Probe.

BPCSS0113I The probe is stopped.

Explanation

No metadata will be collected about the device until the probe is restarted.

Action

Restart the probe, as needed. To start a probe, complete these steps:

1. From the menu bar, go to the resource list page for a resource type. For example, to start a probe for a block storage system, select Resources > Block Storage Systems.
2. Right-click the resource and select Data Collection > Start Probe.

BPCSS0114I The probe is starting at an interval of *interval interval units*.

Explanation

The probe will collect status, configuration, and capacity metadata about the device.

BPCSS0115I The probe interval was updated to *interval*. This change was requested by *user name*.

Explanation

The probe will collect status, configuration, and capacity metadata about the device at the new interval.

BPCSS0116I The probe interval was updated to *interval*.

Explanation

The probe will collect status, configuration, and capacity metadata about the device at the new interval.

BPCSS0117I The probe is enabled. This action was requested by *user name*.

Explanation

The probe will collect status, configuration, and capacity metadata about the device.

BPCSS0118I The probe is enabled.

Explanation

The probe will collect status, configuration, and capacity metadata about the device.

BPCSS0119I The probe is disabled. This action was requested by *user name*.

Explanation

Status, configuration, and capacity metadata will not be collected about the device.

BPCSS0120I The probe is disabled.

Explanation

Status, configuration, and capacity metadata will not be collected about the device.

BPCSS0123I The Probe collection interval is enabled and updated to *interval*. This action was requested by *user name*.

Explanation

The probe will collect status, configuration, and capacity metadata about the device at the new interval.

BPCSS0124I Probe interval is enabled and updated to *interval*.

Explanation

The probe will collect status, configuration, and capacity metadata about the device at the new interval.

BPCSS0105I Probe is starting at an interval of *interval interval units*. This action was requested by *user name*.

Explanation

The probe will collect status, configuration, and capacity metadata about the device at the new interval.

BPCUI - User Interface messages

- [BPCUI0000E The action can't be completed because the following error occurred: Error message text.](#)
- [BPCUI0001E An action could not be completed and the following error message was generated: TPCRemoteException message](#)
- [BPCUI0002E Failed to retrieve the requested data because the service is unavailable.](#)
- [BPCUI0003E The NAPI with the IP address Napi IP was not added because of an Internal Error](#)
- [BPCUI0004E The SSH private key for the NAPI Napi IP could not be uploaded](#)
- [BPCUI0005E The action cannot be completed because the following internal error has occurred: message.](#)
- [BPCUI0007E The discovery job failed to complete.](#)
- [BPCUI0009E The SSH key could not be loaded for the following reason:IOException message](#)
- [BPCUI0010E The host name or IP address that you entered is a resource_type, but you selected to add a different type of storage system.](#)
- [BPCUI0011E The Device Server did not discover any device](#)
- [BPCUI0012E Cannot connect to the device with the address Ip Address.](#)
- [BPCUI0019E No data is available for this selection.](#)
- [BPCUI0025E Probe job job Id failed.](#)
- [BPCUI0029E Invalid parameter param passed.](#)
- [BPCUI0030I This task was already executed.](#)
- [BPCUI0032E An unexpected response was received from the server.](#)
- [BPCUI0034E Invalid number of runs to keep for each schedule. The number should be between param1 and param2.](#)
- [BPCUI0035E Invalid number of days' worth of log-files to keep. The number should be between param1 and param2.](#)
- [BPCUI0036E The schedule id scheduleID associated with this job is no longer valid. It might have been deleted. Refresh the view and try again.](#)
- [BPCUI0037E The replication server is not installed or is unavailable.](#)
- [BPCUI0038E Invalid number of days to retain alerts. The number should be between param1 and param2.](#)
- [BPCUI0039E A Storage Resource agent cannot be found.](#)
- [BPCUI0040E Parsing results from a call to the Data server failed with the following error message: param1.](#)
- [BPCUI0042E Communication with the Data Server failed with the following error: param1](#)
- [BPCUI0043E Cannot connect to the Data server.](#)
- [BPCUI0044E The entity was not found in the database.](#)
- [BPCUI0045E Host name length exceeds the 255 character limit](#)
- [BPCUI0046E Report 'configurationId' not found](#)
- [BPCUI0047E Parameter 'parameterName' is not defined in report configurationId'](#)
- [BPCUI0048E No property is not defined for report configurationId'](#)
- [BPCUI0049E No such property.propertyName for report configurationId'](#)
- [BPCUI0050E variableName can not be overridden](#)
- [BPCUI0051E variableName not valid report output format.](#)
- [BPCUI0052E variableName not reachable](#)
- [BPCUI0053E Cannot authenticate with the provided user credentials.](#)
- [BPCUI0054E The host name or IP address {0} is not valid.](#)
- [BPCUI0055E Cannot connect to the storage system.](#)
- [BPCUI0056E Cannot connect to the storage system or cluster.](#)
- [BPCUI0058I No supported resources were discovered on the data source data_Source_Address.](#)
- [BPCUI0060I File param was successfully uploaded to the Data Server.](#)
- [BPCUI0061E Upload file type param is not supported.](#)
- [BPCUI0062E The requested action failed with the following error message: error message](#)
- [BPCUI0063E Cannot find jobs for scheduleId param and deviceId param. No logs are displayed.](#)
- [BPCUI0064E A log file cannot be displayed for the job.](#)
- [BPCUI0065E The job log file cannot be accessed. The log file may have been manually removed or may have been deleted because it was older than retain_days days or it exceeded the maximum number of no_of_lofs runs.](#)
- [BPCUI0067E The schedule for collecting status and asset data cannot be created.](#)
- [BPCUI0068E A proposed schedule for collecting status and asset data cannot be created.](#)
- [BPCUI0069E The proposed schedule for collecting status and asset data cannot be deleted.](#)
- [BPCUI0071E The task task_name could not be completed.](#)
- [BPCUI0072E Cannot connect to the Device server. Verify that the database service and Device server are running, and that the Device server is accessible.](#)
- [BPCUI0073E Can't make a connection to the storage_resource storage resource.](#)
- [BPCUI0074E The wizard could not set an attribute for the storage resource.](#)

- [BPCUI0075E The certificate wasn't saved on the server.](#)
- [BPCUI0076W The initial job to collect status and asset data did not start.](#)
- [BPCUI0077E A failure occurred loading the certificate.](#)
- [BPCUI0078I The certificate was loaded successfully.](#)
- [BPCUI0079E The SSL certificate is not in the expected format.](#)
- [BPCUI0084W The wizard could not retrieve the default interval information for performance monitoring.](#)
- [BPCUI0085E The user name or password for the hypervisor or vCenter hypervisor or vCenter Server is invalid.](#)
- [BPCUI0086E The SSL certificate is invalid for the hypervisor or vCenter hypervisor or vCenter Server, or the firewall is blocking access to it.](#)
- [BPCUI0087E The version of the hypervisor or vCenter hypervisor or vCenter Server is not supported.](#)
- [BPCUI0088E The host name, protocol, or port for the hypervisor or vCenter hypervisor or vCenter Server is invalid, or the hypervisor or vCenter Server is unreachable.](#)
- [BPCUI0089W Cannot retrieve a valid set of data collection intervals for performance monitoring.](#)
- [BPCUI0090I All alerts were removed.](#)
- [BPCUI0091W error_count of total_count alerts were not removed.](#)
- [BPCUI0093I No data path is available for deviceNameVariable.](#)
- [BPCUI0094E Authorization failed due to an internal error.](#)
- [BPCUI0097E Authorization failed due to an invalid request context.](#)
- [BPCUI0098E The current user is not authorized to perform the requested function.](#)
- [BPCUI0099E The storage resource is not available.](#)
- [BPCUI0100I success_count alerts were marked as acknowledged.](#)
- [BPCUI0101I The alert was marked as acknowledged.](#)
- [BPCUI0102E None of the alerts were marked as acknowledged.](#)
- [BPCUI0104I success_count alerts were marked as unacknowledged.](#)
- [BPCUI0105I The alert was marked as unacknowledged.](#)
- [BPCUI0108I All informational alerts were marked as acknowledged.](#)
- [BPCUI0110W Some informational alerts were not marked as acknowledged.](#)
- [BPCUI0111I All alerts were marked as acknowledged.](#)
- [BPCUI0112I success_count alerts were removed.](#)
- [BPCUI0113I The alert was removed.](#)
- [BPCUI0114I All acknowledged alerts were removed.](#)
- [BPCUI0116W Some acknowledged alerts were not removed.](#)
- [BPCUI0120W Some acknowledged alerts were not marked as unacknowledged.](#)
- [BPCUI0121E Unable to communicate with the product server. Make sure that the server is running properly.](#)
- [BPCUI0122E No job log file was created for this job run.](#)
- [BPCUI0123E The action cannot be completed.](#)
- [BPCUI0124E An unexpected error occurred during the execution of the action.](#)
- [BPCUI0125E The alert is not available.](#)
- [BPCUI0126E The status of the Performance Monitors could not be retrieved.](#)
- [BPCUI0127E The currently installed version of the product does not have the required product license for the function that you requested.](#)
- [BPCUI0128E An undefined capacity chart metric was requested.](#)
- [BPCUI0129I Alerts that were migrated from a previous version of the product are not shown on this page.](#)
- [BPCUI0130E The alerts cannot be acknowledged because they were deleted.](#)
- [BPCUI0131E The alerts cannot be unacknowledged because they were deleted.](#)
- [BPCUI0132W success_count alerts were marked as acknowledged, unsuccess_count alerts cannot be marked as acknowledged because they were deleted.](#)
- [BPCUI0133W success_count alerts were marked as unacknowledged, unsuccess_count alerts cannot be marked as unacknowledged because they were deleted.](#)
- [BPCUI0134E The alert cannot be acknowledged because it was deleted.](#)
- [BPCUI0135E The alert cannot be unacknowledged because it was deleted.](#)
- [BPCUI0136E The device was not removed because the action is not supported for devices of type devType.](#)
- [BPCUI0137E Input text provided has invalid character\(s\): characters. Input text: text](#)
- [BPCUI0141E Host name or IP address hostname specified on line line of file file is not valid.](#)
- [BPCUI0143E Host port WWPn wwpn specified on line line of file file is not valid.](#)
- [BPCUI0144E Duplicate server name specified on lines line1 and line2 of file file.](#)
- [BPCUI0145E Could not parse file file.](#)
- [BPCUI0146E Could not parse file file. Invalid entry on line line.](#)
- [BPCUI0148I Successfully deleted server server_name.](#)
- [BPCUI0149I Successfully modified ports of server server_name.](#)
- [BPCUI0150I The server was created.](#)
- [BPCUI0151E The host name or IP address is associated with another resource.](#)
- [BPCUI0152I The data source data_Source_Address was successfully added as a data source for monitoring. The following new resources were detected:](#)
- [BPCUI0155W You cannot provision volumes because there is no Fibre Channel host port information for at least one server.](#)
- [BPCUI0156W You cannot provision volumes to servers that use different operating systems.](#)
- [BPCUI0157W You cannot provision volumes to servers and virtual machines at the same time. To provision volumes, ensure that you select either only servers or only virtual machines.](#)
- [BPCUI0158I Volumes are assigned to the hypervisors that host virtual machines. Volumes are not assigned directly to virtual machines.](#)
- [BPCUI0159W You cannot provision volumes because at least one of the hypervisors that host the virtual machines is not being monitored. Ensure that all the hypervisors that are hosting the virtual machines that were selected for provisioning were probed.](#)
- [BPCUI0160E Duplicate port WWPn wwpn specified on lines line1 and line2 of file file.](#)
- [BPCUI0162W File file does not contain any servers to create.](#)
- [BPCUI0166W Optimization cannot be done in place to the subsystem since storage subsystem param1 and/or its pools belong to more than one capacity pool. Following are capacity pools the subsystem is associated with: param2](#)
- [BPCUI0167W Optimization cannot be done in place to the subsystem since storage subsystem param1 and/or its pools are not part of any capacity pool.](#)
- [BPCUI0168W Optimization cannot be done in place to the server param1 since storage subsystems or storage pools associated with luns assigned to the server belong to more than one capacity pool. Following are associated capacity pools: param2](#)
- [BPCUI0169W Optimization cannot be done in place to the server param1 since storage subsystems or storage pools associated with luns assigned to the server are not part of any capacity pool.](#)
- [BPCUI0170W Optimization cannot be done in place to the storage entity param1 since storage subsystems or storage pools associated with it belong to more than one capacity pool. Following are associated capacity pools: param2](#)
- [BPCUI0171W Optimization cannot be done in place to the storage entity param1 since storage subsystems or storage pools associated with it are not part of any capacity pool.](#)

- [BPCUI0172E The operation timed out while waiting for a response from the server.](#)
- [BPCUI0173E File does not exist or is empty.](#)
- [BPCUI0174E The device does not support the credential mechanism used.](#)
- [BPCUI0175E A required parameter is missing.](#)
- [BPCUI0176E The highlighted field contains an invalid value.](#)
- [BPCUI0177E The highlighted field contains a value that is outside of the allowed range. The value must be between minVal and maxVal.](#)
- [BPCUI0178E A service class with the same name and type already exists.](#)
- [BPCUI0179I The service class was created.](#)
- [BPCUI0180I Based on the known configuration of storage system host connections, fabric zone aliases, and HBA ports, additional ports may have been added to the selection below.](#)
- [BPCUI0181I You selected to add a expectedDevice resource, but a foundDevice resource was detected and will be added.](#)
- [BPCUI0182I The data source data_Source_Address was added as a data source for monitoring. No new resources were detected.](#)
- [BPCUI0183E The text in the highlighted field exceeds the maxLength character limit.](#)
- [BPCUI0185W Unable to lookup the IP Address for Host Name hostName. Enter the IP Address manually.](#)
- [BPCUI0189I Configuration of SRA deployment and probe schedules were done successfully.](#)
- [BPCUI0190W Configuration of SRA finished with some warnings or errors. Check the detail messages.](#)
- [BPCUI0191E An internal error occurred while testing connection to param1.](#)
- [BPCUI0192E The supplied service class type is invalid.](#)
- [BPCUI0193E The specified SMI-S provider was not found. Make sure that the protocol, SMI-S provider host name or IP address, and port are specified correctly and that the SMI-S provider is properly configured at that location.](#)
- [BPCUI0194E An unknown error has occurred. Please review all values entered.](#)
- [BPCUI0195E The Interop Namespace is not correct. Please correct this entry.](#)
- [BPCUI0196E A timeout occurred while processing the request. Please retry request.](#)
- [BPCUI0197E A connection was not established. Make sure that the protocol, SMI-S provider host name or IP address, and port are specified correctly.](#)
- [BPCUI0198E The authentication to the SMI-S provider failed.](#)
- [BPCUI0199E An SSLHandshakeException or SSLProtocolException has occurred. This exception might be due to an invalid SLP registration, e.g., 'http' instead of 'https'.](#)
- [BPCUI0201E There is a pending delete in process for this SMI-S provider.](#)
- [BPCUI0202I Success](#)
- [BPCUI0203E The selected resources were not removed.](#)
- [BPCUI0204W successfulDeletes of attemptedDeletes of the selected resources were removed.](#)
- [BPCUI0205W successfulDeletes selected resources were removed, however warnings did occur.](#)
- [BPCUI0209E A database operation cannot be completed.](#)
- [BPCUI0210I Device param1 supports performance monitoring.](#)
- [BPCUI0211E No performance data is available for a resource.](#)
- [BPCUI0212E There is no Secure Shell running at this host/IP.](#)
- [BPCUI0213E Unsupported Secure Shell protocol was used.](#)
- [BPCUI0214E Invalid public key location for subsystem param1.](#)
- [BPCUI0215E Invalid public key format for subsystem param1.](#)
- [BPCUI0216E Passphrase was incorrect for subsystem param1.](#)
- [BPCUI0217E Unable to transfer the key\(s\) to the server param1.](#)
- [BPCUI0218E The specified private key file format is not supported. Please convert it to Open SSH \(.pem\) key file format for subsystem param1.](#)
- [BPCUI0219E The specified key file or key file name is already linked to another user.](#)
- [BPCUI0220E The IP address that was entered was the address of the management console for the storage system. You must enter the valid IP address of the block component of the storage system.](#)
- [BPCUI0221E The IP address you entered is the address of another device's management console.](#)
- [BPCUI0222E The IP address you entered points to a device of another type.](#)
- [BPCUI0223E Passphrase is required. Specify one for subsystem param1.](#)
- [BPCUI0224E Cannot connect to a resource because of an SSL certificate error. Troubleshooting information: http://www.ibm.com/support/docview.wss?uid=swg21976237](#)
- [BPCUI0225I The agent log files for server_Name have been collected and copied to log_Location.](#)
- [BPCUI0226I Discovery of data_source is taking longer than expected. Click Close to run the discovery in the background.](#)
- [BPCUI0227E Thin provisioning must be enabled when compression is enabled.](#)
- [BPCUI0229I 1 resource was added to name.](#)
- [BPCUI0231I count resources were added to name.](#)
- [BPCUI0233E The specified host name is already associated with an existing server.](#)
- [BPCUI0234E The specified IP address is already associated with an existing server.](#)
- [BPCUI0235E The specified host name and IP address are already associated with an existing server.](#)
- [BPCUI0236E The disabling of the agents failed.](#)
- [BPCUI0237E Errors occurred when attempting to disable some of the agents.](#)
- [BPCUI0238W Warnings occurred when attempting to disable warningCount of the agents.](#)
- [BPCUI0239I attemptedCount of the selectedCount selected agents were disabled.](#)
- [BPCUI0240E The agents were not enabled.](#)
- [BPCUI0241E Errors occurred when attempting to enable some of the agents.](#)
- [BPCUI0242W Warnings occurred when attempting to enable warningCount of the agents.](#)
- [BPCUI0243I attemptedCount of the selectedCount selected agents were enabled.](#)
- [BPCUI0244I The credentials of an agent were updated.](#)
- [BPCUI0245I The credentials of updateCount agents were updated.](#)
- [BPCUI0246E Cannot authenticate to the file module with the provided user credentials.](#)
- [BPCUI0247E Unknown file module key user.](#)
- [BPCUI0248E The SSH key could not be loaded for the following reason:IOException message](#)
- [BPCUI0249E Passphrase is incorrect.](#)
- [BPCUI0250E Passphrase is required.](#)
- [BPCUI0251E Cannot connect to the storage system or cluster.](#)
- [BPCUI0252E The host name or IP address {Q} is not valid.](#)
- [BPCUI0253E Cannot connect to the data source for the resource with the address ip_address.](#)
- [BPCUI0254E Invalid private key location.](#)
- [BPCUI0255W The following resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?](#)
- [BPCUI0256W The following resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?](#)
- [BPCUI0257W The following resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?](#)

- [BPCUI0258W The following internal resources of a storage system you are attempting to add are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?](#)
- [BPCUI0259W The following storage systems and storage-system internal resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?](#)
- [BPCUI0260E The specified private key file format for the file module is not supported. Please convert it to Open SSH \(.pem\) key file.](#)
- [BPCUI0261E The service class was not found in the database.](#)
- [BPCUI0262E The capacity pool was not found in the database.](#)
- [BPCUI0263E The scheduling of the agent upgrade jobs failed.](#)
- [BPCUI0264E Errors occurred when attempting to schedule the upgrade jobs of some of the agents.](#)
- [BPCUI0265W Warnings occurred when scheduling the upgrade of warningCount of the agents.](#)
- [BPCUI0266I AttemptedCount of the selected agents were scheduled for upgrade.](#)
- [BPCUI0267I The upgrade agent job was successfully scheduled for hostName.](#)
- [BPCUI0268W Deleting a capacity pool does not affect any volumes or shares that were provisioned from the capacity pool. However, the volumes or shares are no longer associated with the capacity pool. Associations with the following volumes or shares will be removed:](#)
- [BPCUI0269W The following volumes are associated with the service class scName. When the volumes were created, they satisfied the requirements of the service class. If you modify the service class, the volumes are still associated with the service class, but might not satisfy the new requirements of the service class. Depending on your changes to the service class, users might incorrectly assume that the volumes have properties that they do not possess.](#)
- [BPCUI0270W The following shares are associated with the service class scName. When the shares were created, they satisfied the requirements of the service class. If you modify the service class, the shares are still associated with the service class, but might not satisfy the new requirements of the service class. Depending on your changes to the service class, users might incorrectly assume that the shares have properties that they do not possess.](#)
- [BPCUI0271W The following volumes are associated with the service class scName. If you delete the service class, the volumes are no longer associated with any service class.](#)
- [BPCUI0272W The following shares are associated with the service class scName. If you delete the service class, the shares are no longer associated with any service class.](#)
- [BPCUI0273E The action does not support the specified type of device.](#)
- [BPCUI0274I The connection test to resource data_Source_Name was successful.](#)
- [BPCUI0275I To collect data about zoning or complete zoning actions during provisioning, you must deploy Storage Resource agents to one or more servers that are on the fabric.](#)
- [BPCUI0276I Agent agentName was disabled.](#)
- [BPCUI0277I Agent agentName was enabled.](#)
- [BPCUI0278I The credentials for agentName were updated.](#)
- [BPCUI0279I There is no job defined for the device Name. Please create a job first before running it again.](#)
- [BPCUI0280I No switches are managed by the data_Source_Address data source.](#)
- [BPCUI0282I The resources that are managed by data_Source_Address are already known. One or more resources were added.](#)
- [BPCUI0284I No fabrics are managed by the data_Source_Address data source.](#)
- [BPCUI0286I The fabrics that are managed by data_Source_Address are already being monitored.](#)
- [BPCUI0289W The following network shared disks \(NSDs\) are already assigned to a capacity pool. Are you sure you want to move these NSDs to a different capacity pool?](#)
- [BPCUI0290W The following file systems and network shared disks \(NSDs\) are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?](#)
- [BPCUI0291W The following network shared disks \(NSDs\) are already assigned to a capacity pool. Are you sure you want to move these NSDs to a different capacity pool?](#)
- [BPCUI0292E The host name or IP address ip_address_or_hostname cannot be reached.](#)
- [BPCUI0293I A probe is started for deviceName.](#)
- [BPCUI0294I A performance monitor is started for deviceName.](#)
- [BPCUI0295I The performance monitor is stopped for deviceName.](#)
- [BPCUI0297W One resource was added to capacity_pool_name. One resource could not be added because it could not be found.](#)
- [BPCUI0298W count resources were added to capacity_pool_name. One resource could not be added because it could not be found.](#)
- [BPCUI0299W One resource was added to capacity_pool_name. count_Not_Found resources could not be added because they could not be found.](#)
- [BPCUI0300W count resources were added to capacity_pool_name. count_Not_Found resources could not be added because they could not be found.](#)
- [BPCUI0301E Failed to assign the role name role.](#)
- [BPCUI0302E Failed to retrieve the existing role assignments.](#)
- [BPCUI0303E Failed to remove all role assignments from the specified groups.](#)
- [BPCUI0304W An error occurred when saving the user-defined properties of the resourcesType.](#)
- [BPCUI0305E A capacity pool with the same name already exists.](#)
- [BPCUI0306W The selected resource was removed, however warnings did occur.](#)
- [BPCUI0307E The schedule could not be deleted.](#)
- [BPCUI0308I The resource does not have a connection configured. To add a connection to the resource, click Add Storage System.](#)
- [BPCUI0309I A probe schedule is defined for deviceName.](#)
- [BPCUI0310I A performance monitor schedule is defined for deviceName.](#)
- [BPCUI0311I Probe and performance monitor schedules are defined for deviceName.](#)
- [BPCUI0312I SNMP Discovery of switches is taking longer than expected. Click Close to run the discovery in the background.](#)
- [BPCUI0313I An upgrade is started for server deviceName.](#)
- [BPCUI0314E Failed to retrieve the list of user groups from the WebSphere user repository.](#)
- [BPCUI0315E Failed to retrieve the list of user groups from user repository due to an invalid search string.](#)
- [BPCUI0316W Failed to update the role cache maintained by the Device server.](#)
- [BPCUI0317E Access can not be removed, because at least one Administrator user must remain in the system.](#)
- [BPCUI0318E The group mapping can not be modified, because at least one Administrator user must remain in the system.](#)
- [BPCUI0319I A task is started for resource resourceName.](#)
- [BPCUI0320I Probe and performance monitor schedules are defined for deviceName. A performance monitor is scheduled to collect performance data after the probe is done.](#)
- [BPCUI0321I A task is paused for resource resourceName.](#)
- [BPCUI0322E A task could not be paused for resource resourceName.](#)
- [BPCUI0323I A task is resumed for resource resourceName.](#)
- [BPCUI0324E A task could not be resumed for resource resourceName.](#)
- [BPCUI0325E Failed to retrieve the list of users from the WebSphere user repository.](#)
- [BPCUI0326E Failed to retrieve the list of users from user repository due to an invalid search string.](#)
- [BPCUI0327E Failed to get the roles associated with the current user.](#)
- [BPCUI0328I A task is saved.](#)
- [BPCUI0329I A task was successfully deleted.](#)

- [BPCUI0330E The user user is not authorized to access the product.](#)
- [BPCUI0331I A task is cancelled for resource resourceName.](#)
- [BPCUI0332E An unexpected error occurred. The task for schedule schedule name could not be paused or resumed.](#)
- [BPCUI0333E An unexpected error occurred. The task for schedule schedule name could not be paused.](#)
- [BPCUI0334E An unexpected error occurred. The task for schedule schedule name could not be resumed.](#)
- [BPCUI0335E The volumes cannot be converted or moved because the target pools do not have sufficient available space.](#)
- [BPCUI0336I The ability to provision with block storage devices is only available with the advanced license.](#)
- [BPCUI0338E Insufficient user privileges to service the REST request.](#)
- [BPCUI0339E An unexpected error occurred while authorizing the REST request.](#)
- [BPCUI0340I A task was successfully renamed.](#)
- [BPCUI0341E The task could not be renamed.](#)
- [BPCUI0342E The task could not be renamed because the specified name already exists.](#)
- [BPCUI0343I Performance monitoring is unavailable for resource resource name because the resource was not probed.](#)
- [BPCUI0344W The following service classes allow provisioning only from the capacity pool capacity pool: service classes. If you delete this capacity pool, the service classes will allow provisioning from any available storage.](#)
- [BPCUI0346I The Storage Resource agent that is deployed on the server cannot be uninstalled.](#)
- [BPCUI0347I All servers were removed except for the product server. Entries for the product server resources might still be displayed in the GUI until all the associated removals are complete.](#)
- [BPCUI0348W You cannot provision volumes because at least one of the selected hosts was not found in the database. Ensure that all hosts that are selected for provisioning are being monitored.](#)
- [BPCUI0349W You cannot provision volumes because not all of the selected hosts appear to have Fibre Channel connectivity.](#)
- [BPCUI0350W You cannot provision volumes because the hypervisors that host the virtual machines use different operating systems.](#)
- [BPCUI0351W You cannot provision volumes because there is no Fibre Channel host port information for at least one hypervisor.](#)
- [BPCUI0352W You cannot provision volumes because not all of the hypervisors that host the virtual machines appear to have Fibre Channel connectivity.](#)
- [BPCUI0355W You cannot provision volumes because no block-storage service class exists.](#)
- [BPCUI0356W You cannot provision shares because no file-storage service class exists.](#)
- [BPCUI0357W You cannot provision volumes because you do not have permission to provision by using any block-storage service class.](#)
- [BPCUI0358W You cannot provision shares because you do not have permission to provision by using any file-storage service class.](#)
- [BPCUI0359E The credentials for the servers were not updated.](#)
- [BPCUI0360W The credentials for successfulUpdates of attemptedUpdates of the selected servers were updated.](#)
- [BPCUI0361W The credentials for the selected server was updated, however warnings did occur.](#)
- [BPCUI0362W The credentials for successfulUpdates selected servers were updated, however warnings did occur.](#)
- [BPCUI0363E Cannot connect to the SNMP data source IP_Address.](#)
- [BPCUI0364I The performance monitor schedule was updated for deviceName.](#)
- [BPCUI0366W The server serverName was not updated because it does not support the action.](#)
- [BPCUI0367W You cannot provision volumes to virtual machines with NPIV ports and virtual machines without NPIV ports at the same time. To provision volumes to virtual machines, ensure that you select either only virtual machines with NPIV ports or only virtual machines without NPIV ports.](#)
- [BPCUI0368W You cannot provision volumes because none of the selected hosts appear to have Fibre Channel connectivity and the automatic zoning option is enabled. Disable the automatic zoning option in your zoning policy.](#)
- [BPCUI0369W You cannot provision volumes because none of the hypervisors that manage the selected virtual machines appear to have Fibre Channel connectivity and the automatic zoning option is enabled. Disable the automatic zoning option in your zoning policy.](#)
- [BPCUI0370E The display name displayName is already assigned to resource resource Name.](#)
- [BPCUI0372I The selected hosts do not appear to have Fibre Channel connectivity. In the resulting provisioning task, ensure that the recommended storage system is connected to the hosts before you run the task. Also, be aware that all fabric-related options will be ignored.](#)
- [BPCUI0373I Volumes are assigned to the hypervisors that host virtual machines. Volumes are not assigned directly to virtual machines that do not have NPIV ports. None of the hypervisors that manage the virtual machines appear to have Fibre Channel connectivity. In the resulting provisioning task, ensure that the recommended storage system is connected to the hypervisors before you run the task. Also, be aware that all fabric-related options will be ignored.](#)
- [BPCUI0374E Schedule is not enabled for the resource resource.](#)
- [BPCUI0375E Performance data is not available.](#)
- [BPCUI0376E Invalid number of days to keep configuration history. The number should be between minimum value and maximum value.](#)
- [BPCUI0377E Invalid number of days to keep data for removed resources. The number should be between minimum value and maximum value.](#)
- [BPCUI0378E Invalid number of days to keep sample performance data. The number should be between minimum value and maximum value.](#)
- [BPCUI0379E Invalid number of days to keep hourly performance data. The number should be between minimum value and maximum value.](#)
- [BPCUI0380E Invalid number of days to keep daily performance data. The number should be between minimum value and maximum value.](#)
- [BPCUI0381E Failed to update the performance data retention settings.](#)
- [BPCUI0382E Performance monitoring is unavailable for resource resource name.](#)
- [BPCUI0383E Failed to update the history retention settings.](#)
- [BPCUI0384E Failed to retrieve the history retention settings.](#)
- [BPCUI0385E Invalid number of runs to keep log files for each schedule. The number should be between minimum value and maximum value.](#)
- [BPCUI0386E A job cannot be run for resource resourceName because there is a job already running for the resource. Wait for the job to finish and try again.](#)
- [BPCUI0387I The selected resources support different performance monitor intervals. If you select multiple resources, intervals that are common to all resources are displayed in the interval list.](#)
- [BPCUI0388E The probe schedule cannot be created for resource {0} because not all the information was provided. If you are configuring a probe for a resource for the first time, you must enter values for the probe status, time, and frequency fields.](#)
- [BPCUI0389E The performance monitor schedule cannot be created because not all the information was provided. If you are configuring a performance monitor for a resource for the first time, you must enter values for the performance monitor status and interval fields.](#)
- [BPCUI0390I The service logs were successfully created.](#)
- [BPCUI0391I The connection test to data source data source was successful. A probe is running. The health status is unknown until the probe is finished.](#)
- [BPCUI0392I The connection test to the data source data source was successful.](#)
- [BPCUI0393E The user user_name does not have sufficient privileges to deploy the vSphere Web Client extensionr.](#)
- [BPCUI0394E The user user_name does not have permission to log in to the vCenter Server system.](#)
- [BPCUI0395E This version of the vCenter Server server_name does not support the deployment of the vSphere Web Client extension for the product.](#)
- [BPCUI0396E The user user_ID does not have the required role. The role associated with this user must be Administrator, Monitor, or External Application.](#)
- [BPCUI0397E The vCenter Server user name or password is invalid.](#)
- [BPCUI0398E The user name or password is invalid.](#)
- [BPCUI0399I The server was started.](#)
- [BPCUI0400E Failed to retrieve the system management information from the Data server.](#)
- [BPCUI0402E Failed to retrieve the server status of the Data server.](#)
- [BPCUI0403E The SMI-S provider service is not available.](#)

- [BPCUI0404E An error occurred while updating the trace log configuration file. The original file file was deleted and could not be restored. A backup of this file may be available at backup file.](#)
- [BPCUI0405E Failed to set the trace settings from the Data server.](#)
- [BPCUI0406E Cannot start the server. The start script reported the following error: error](#)
- [BPCUI0407E Cannot start the server. Unable to locate the start script path to script.](#)
- [BPCUI0408E Cannot start the server. Unable to execute the start script path to script.](#)
- [BPCUI0409W The server is taking a long time to start. If the server status continues to show an error status after a reasonable interval, try to start the server again.](#)
- [BPCUI0410E Cannot stop the server. The stop script reported the following error: error](#)
- [BPCUI0411W The server is taking a long time to stop. If the server status continues to show that it is still running try to stop the server again after a reasonable interval.](#)
- [BPCUI0412E Cannot stop the server. Unable to locate the stop script path to script.](#)
- [BPCUI0413E Cannot stop the server. Unable to execute the stop script path to script.](#)
- [BPCUI0414W It is taking a long time for the services to start. If the server status continues to show an error status after a reasonable interval, try to start the services again. If the problem persists then restart the server.](#)
- [BPCUI0415E Failed to start the service service name.](#)
- [BPCUI0416I The server was stopped.](#)
- [BPCUI0417I The services of the server were started.](#)
- [BPCUI0418E The action cannot be completed because the data source that is managing this resource cannot be reached.](#)
- [BPCUI0419E A Storage Resource agent is already deployed for this server and has a status of Pending deployment or Failed deployment. Use the Servers page to resolve the deployment errors or modify the deployment schedule.](#)
- [BPCUI0420E A file access error occurred when the system attempted to back up or modify the tracing configuration file configuration file.](#)
- [BPCUI0421E There is a log collection operation already running. A new one cannot be submitted until the current one completes.](#)
- [BPCUI0422E Cannot start the log collecting job. Unable to locate the required script path to script.](#)
- [BPCUI0423E Cannot start the log collecting job. Unable to run the log collection script path to script.](#)
- [BPCUI0424E Storage cannot be provisioned from capacity pool capacity pool using service class service class for the following reason:](#)
- [BPCUI0425W The task task name cannot be scheduled because it is already running.](#)
- [BPCUI0426E Storage cannot be provisioned by using service class service class for the following reason:](#)
- [BPCUI0427W The selected group action is complete for all tasks, but warnings were reported.](#)
- [BPCUI0428I The selected group action is complete for all tasks. Some informational messages were returned.](#)
- [BPCUI0429E The validation process cannot contact the server. The server might be down or unreachable due to network problems.](#)
- [BPCUI0430I Some tasks were not deleted because they were already run.](#)
- [BPCUI0431E Failed to retrieve the list of managed devices.](#)
- [BPCUI0432E Failed to retrieve the performance monitoring granularity from the Device server. Check the connection to the Device server and retry the operation.](#)
- [BPCUI0433E OS type osType specified on line line of file file is not valid.](#)
- [BPCUI0434E Data source data_Source_Key could not be found.](#)
- [BPCUI0435E Required host name or IP address and OS type were not specified on line line of file file.](#)
- [BPCUI0436E The alert notification settings cannot be displayed.](#)
- [BPCUI0437E The alert notification settings cannot be saved.](#)
- [BPCUI0438E File file does not exist or is empty.](#)
- [BPCUI0439E The file file could not be uploaded.](#)
- [BPCUI0440E The text location specified on line line of file file has invalid character\(s\): characters](#)
- [BPCUI0441E The alert definitions cannot be displayed.](#)
- [BPCUI0442E The alert definitions cannot be saved.](#)
- [BPCUI0443E Select at least one managed server that is deployed for which alert notification settings need to be displayed.](#)
- [BPCUI0444E Select at least one managed server that is deployed for which alert definitions need to be displayed.](#)
- [BPCUI0445W The discovery job completed with errors. Some available devices were not discovered.](#)
- [BPCUI0446E Unable to test the connection to the device because the request was not processed by the data collector.](#)
- [BPCUI0447E Select at least one managed storage subsystem for which alert notification settings need to be displayed.](#)
- [BPCUI0448E Select at least one managed storage subsystem for which alert definitions need to be displayed.](#)
- [BPCUI0449E The user does not have the required authority to complete the task or command.](#)
- [BPCUI0451E One or more applications from provided list: names do not exist.](#)
- [BPCUI0452E entity name is not supporting data collection actions.](#)
- [BPCUI0453E One or more departments from provided list: names do not exist.](#)
- [BPCUI0455I No performance data is available for the selected resources.](#)
- [BPCUI0456E You cannot complete the action because the service is temporarily unavailable.](#)
- [BPCUI0457W The applications listOfApplications cannot be deleted because they contain subcomponents subcomponent, which cannot be moved up a level in the applications hierarchy due to name conflicts with existing applications in that higher level.](#)
- [BPCUI0458W The departments listOfDepartments cannot be deleted because they contain subdepartments or applications subdepartment, which cannot be moved up a level in the departments hierarchy due to name conflicts with departments in that higher level.](#)
- [BPCUI0459W The selected subcomponents cannot be removed from the application because they cannot be moved up a level in the application hierarchy due to name conflicts with the existing applications or subcomponents at the higher level.](#)
- [BPCUI0460W The selected applications or subdepartments cannot be removed from the department because they cannot be moved up a level in the department hierarchy due to name conflicts with the existing applications or subdepartments at the higher level.](#)
- [BPCUI0461W There are no task details to display. The analysis-execution task could not be run.](#)
- [BPCUI0462E Failed to add the device because the data collector is not responding.](#)
- [BPCUI0463E The discovery failed because the data collector is not responding.](#)
- [BPCUI0464E The connection test failed because the data collector is not responding.](#)
- [BPCUI0465E The requested action failed because the data collector is not responding.](#)
- [BPCUI0466I The servers were created.](#)
- [BPCUI0467W successCount of totalCount servers were created.](#)
- [BPCUI0468E The creation of the servers failed.](#)
- [BPCUI0469E Schedule job does not exist for entity name.](#)
- [BPCUI0470E Invalid file file size of size GB. Maximum allowed file size is max size GB.](#)
- [BPCUI0471E Failed to set the trace settings from the Alert server.](#)
- [BPCUI0472E Failed to retrieve the system management information from the Alert server.](#)
- [BPCUI0474E Failed to retrieve the server status of the Alert server.](#)
- [BPCUI0475I The volumes have been excluded from the reclamation analysis.](#)
- [BPCUI0476I The volumes will be included in future analyses to reclaim storage.](#)

- [BPCUI0477E An unexpected error occurred when modifying the optimization characteristics of the volumes.](#)
- [BPCUI0478E The scheduled agent upgrade time is in the past.](#)
- [BPCUI0479E The object storage credentials are incorrect. Enter the correct credentials. Alternatively, clear the object credentials check box and do not specify the authentication credentials for object storage now. You can use the Modify Connection action to add the object storage later.](#)
- [BPCUI0480E An object storage request failed on the GPFS cluster.](#)
- [BPCUI0481W No resources were removed.](#)
- [BPCUI0482E No resources were updated.](#)
- [BPCUI0483E The connection information cannot be updated because it points to another device.](#)
- [BPCUI0484I The connection information for device name was updated.](#)
- [BPCUI0485E The connection information cannot be updated.](#)
- [BPCUI0486E Cannot query the object service for information about accounts and containers as the specified user does not have admin privileges.](#)
- [BPCUI0487I The connection information of the selected device was successfully updated. Other devices were detected as being managed by the same data source. Would you like to update the connection information of all of them?](#)
- [BPCUI0488I The connection information of all devices connecting through this data source was updated.](#)
- [BPCUI0489W Some of the devices connecting through this data source failed to be updated.](#)
- [BPCUI0490I The vCenter vCenter Server was removed.](#)
- [BPCUI0491E The vCenter vCenter Server was not found in the database.](#)
- [BPCUI0492E The selected vCenter Servers were not found in the database.](#)
- [BPCUI0493I The vCenter vCenter Server and all number of monitored hypervisors hypervisors monitored by it were successfully removed.](#)
- [BPCUI0494I The number of vCenters selected vCenter Servers and all number of monitored hypervisors hypervisors monitored by them were successfully removed.](#)
- [BPCUI0495W Only number of removed vCenters of number of selected vCenters of the selected vCenter Servers and number of removed monitored hypervisors of number of monitored hypervisors of the hypervisors monitored by them were successfully removed.](#)
- [BPCUI0496I The following fabrics were detected as being managed by the same data source: comma separated fabrics list. This action applies to all fabrics that are managed by the current data source. Would you like to update the connection information of all of them?](#)
- [BPCUI0497E The following fabrics cannot be monitored through the SMI agent: comma separated fabrics list. The data source connection information will not be updated.](#)
- [BPCUI0498E The fabric cannot cannot be monitored through the SMI agent.](#)
- [BPCUI0499I Other switches were detected as being managed by the same data source. This action applies to all switches that are managed by the current data source. Would you like to update the connection information of all of them?](#)
- [BPCUI0500E One or more switches cannot be monitored through the SMI agent. The data source connection information will not be updated.](#)
- [BPCUI0501E The information cannot be displayed. Log out of the GUI, log in, and try the action again.](#)
- [BPCUI0502E The device is already managed by this data source. The data source connection information will not be updated.](#)
- [BPCUI0503I The connection information of the selected switches was updated.](#)
- [BPCUI0504I The detected versions of the resources discovered on the data source data_Source_Address are unsupported.](#)
- [BPCUI0505E The resource does not have a connection configured.](#)
- [BPCUI0506E Cannot connect to the Alert server.](#)
- [BPCUI0507E The version of the tpc_server IBM Spectrum Control Server is not supported.](#)
- [BPCUI0508E Cannot connect to the rollup server rollup_server on port host_port.](#)
- [BPCUI0509E Cannot authenticate with the rollup server using the provided credentials.](#)
- [BPCUI0510E You entered an invalid time range. The start date and time must be before the end date and time.](#)
- [BPCUI0511E The following alert name\(s\) are not unique: names.](#)
- [BPCUI0512E Custom alerts already exist for other resources with the following alert name\(s\): names.](#)
- [BPCUI0513E Unable to connect from rollup server rollup_server to the repository database.](#)
- [BPCUI0514E The specified secondary server secondary_server is the primary server.](#)
- [BPCUI0515E The duration of the automated probe run window must be at least minimum_hours hours.](#)
- [BPCUI0516W The selected subgroups cannot be removed from the general group because they cannot be moved up a level in the groups hierarchy due to name conflicts with the general groups at the higher level.](#)
- [BPCUI0519E Authorization has failed because the private key is not valid for the user name that you have specified.](#)
- [BPCUI0520E The IP address ip_address for the FlashSystem storage system is not the management IP address.](#)
- [BPCUI0521E The configuration for the report can't be saved.](#)
- [BPCUI0522E Failed to delete a report configuration.](#)
- [BPCUI0523E Alerts cannot be defined for this storage system.](#)
- [BPCUI0524E The changes to the report configuration can't be saved.](#)
- [BPCUI0525E The configuration for the report can't be saved because the report title isn't unique.](#)
- [BPCUI0527E The action cannot be completed because of an invalid request.](#)
- [BPCUI0528E The action cannot be completed because of an invalid file upload request.](#)
- [BPCUI0526I The connection test to data source data_source was successful. A probe is running.](#)
- [BPCUI0529I The data source data_Source_Address is already being managed as a data source for monitoring. No new resources were detected.](#)
- [BPCUI0530I The data source data_Source_Address is already being managed as a data source for monitoring. The following new resources were detected:](#)
- [BPCUI0531E The action cannot be completed because LDAP registry file failed to upload.](#)
- [BPCUI0532E The action failed because of a missing resource.](#)
- [BPCUI0533E The LDAP configuration test failed.](#)
- [BPCUI0534E There was an error executing the collect log process. If this problem persists, you can try collecting and uploading the service logs manually. Learn More.](#)
- [BPCUI0535E An FTP connection can not be established. If your organization requires the use of a proxy server, consult the following documentation: Troubleshooting FTP Transfers.](#)
- [BPCUI0536E The support data collection failed due to an invalid PMR number format.](#)
- [BPCUI0537E The support package could not be created because file system permissions prevent the creation of temporary files.](#)
- [BPCUI0538E The support data collection completed creating a support package, but the package could not be uploaded to IBM.](#)
- [BPCUI0539E The support data collection failed with an internal error.](#)
- [BPCUI0540E The support data collection failed due to an invalid email address format.](#)
- [BPCUI0541E The specified SMI agent was not found. Make sure that the protocol, SMI agent host name or IP address, and port are specified correctly and that the SMI agent is properly configured at that location.](#)
- [BPCUI0542E A connection was not established. Make sure that the protocol, SMI agent host name or IP address, and port are specified correctly.](#)
- [BPCUI0543E The authentication to the SMI agent failed.](#)
- [BPCUI0544E There is a pending delete in process for this SMI agent.](#)
- [BPCUI0545E The SMI agent service is not available.](#)
- [BPCUI0546E The action cannot be completed because the LDAP registry file could not be updated.](#)
- [BPCUI0547E Connection failed. The server might be down or unreachable due to network problems.](#)

- [BPCUI0548E The add SSL certificate action failed.](#)
- [BPCUI0549E The add SSL certificate action failed because of a wrong password.](#)
- [BPCUI0550E The specified storage resource is not valid for the REST API service request.](#)
- [BPCUI0551E The file cannot be used because it is not a valid SSL certificate. Select a valid certificate file and try again.](#)
- [BPCUI0554E The SSL certificate download process failed.](#)
- [BPCUI0555E The test connection to the LDAP server failed. Verify that your XML file contains the correct syntax and values and that the LDAP server is running.](#)
- [BPCUI0556E An unexpected error occurred creating or updating a support ticket.](#)
- [BPCUI0557E An invalid request was made when creating or updating a support ticket.](#)
- [BPCUI0558E This tier name is already in use. Enter a different name.](#)
- [BPCUI0559E The custom dashboard was removed by another user. Cancel the action and refresh the page manually.](#)
- [BPCUI0600W Can't save the scheduling information for the report because the Data server is offline.](#)
- [BPCUI0601I The resource does not have a connection configured. To add a connection to the resource, click Add Switch or Add Fabric.](#)
- [BPCUI0602E The osAuthentication script does not start. The script reported the following error: script_error.](#)
- [BPCUI0603E The connection test to data source data source was not successful.](#)
- [BPCUI0604E Can't stop data collection for entity name.](#)
- [BPCUI0605E Can't restart data collection for entity name.](#)
- [BPCUI0606E The action cannot be completed because there was a failure to create or write into the pending configuration file.](#)
- [BPCUI0607E The action cannot be completed because there was a failure to read the pending LDAP registry file.](#)
- [BPCUI0608E The action cannot be completed because there was a failure to get the list of LDAP groups.](#)
- [BPCUI0609E The Local OS authentication configuration test failed.](#)
- [BPCUI0610E Failed to update modified username for IBMid unique ID.](#)
- [BPCUI0611E Failed to update IBMid unique ID for IBMid username.](#)
- [BPCUI0612E The action was not performed due to invalid device credentials for entity name.](#)
- [BPCUI0613E A switch with this host name or IP address is already being monitored.](#)

BPCUI0000E The action can't be completed because the following error occurred: *Error message text.*

Explanation

The specified error occurred while processing a request.

Action

To resolve the issue, try the following actions:

- Verify that the local area network is available and a firewall is not preventing network access to product services and agents.
- Check the status of the product servers and database repository on the Home > System Management page.
- On Windows, verify that the related database services are active.
- Check for error messages in the log files for the servers.

If the problem persists, check the product's log files for error messages that might help determine the problem. For information about the location of log files, check the IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93/>.

Related reference

- [IBM Spectrum Control documentation](#)

BPCUI0001E An action could not be completed and the following error message was generated: *TPCRemoteException message*

Explanation

The action could not be completed because of an error related to the Device server component.

Action

Try the following actions:

- Check the status of the Device server on the Home > System Management page.
- Check for error messages in the log file of the Device server.
- Verify that that the database repository is up and running.
- Try the action again.

Related reference

- [Troubleshooting problems with the IBM Spectrum Control component and servers](#)
- [Default locations of log files](#)

BPCUI0002E Failed to retrieve the requested data because the service is unavailable.

Explanation

Data cannot be retrieved from the database repository. This problem might occur if the database repository is unavailable or communication with the database cannot be established.

Action

Verify that the database repository is up and the related database service is active. Ensure that all required servers are running and that the local area network is available. Verify that you have a network connection to the server on which the database repository is located. Check the log files of the servers for error messages that might help determine the problem. See the product information center for the location of these log files.

Related reference

- [🔗 Troubleshooting problems with the IBM Spectrum Control component and servers](#)
- [🔗 Default locations of log files](#)

BPCUI0003E The NAPI with the IP address *Napi IP* was not added because of an Internal Error

Explanation

Action

BPCUI0004E The SSH private key for the NAPI *Napi IP* could not be uploaded

Explanation

Action

BPCUI0005E The action cannot be completed because the following internal error has occurred: *message*.

Explanation

An error occurred while processing a user request.

Action

Wait a few minutes and try again. If you still can't complete the action, go to your Products and services page (<https://myibm.ibm.com/products-services/>) on IBM Marketplace. Click the down-arrow for the Storage Insights offering, click Support, and then choose an option.

Related reference

- [🔗 Getting support](#)
- [🔗 Default locations of log files](#)

BPCUI0007E The discovery job failed to complete.

Explanation

No storage resources were detected because the discovery job did not complete.

Action

Check the Device server log files for error messages that might help determine why the discovery job failed. For the location of the log files, go to IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93/> and view the Reference section.

Related reference

- [IBM Spectrum Control documentation](#)

BPCUI0009E The SSH key could not be loaded for the following reason: *IOException message*

Explanation

Upload of the file failed. Please retry the upload.

BPCUI0010E The host name or IP address that you entered is a *resource_type*, but you selected to add a different type of storage system.

Explanation

The model of the storage system that you add must match the type of storage system that you selected to add.

BPCUI0011E The Device Server did not discover any device

Explanation

The Device Server did not discover any device

BPCUI0012E Cannot connect to the device with the address *Ip Address*.

Explanation

The storage device with this IP address cannot be reached.

Action

Verify that the provided user name, password, and address for the device are correct. Make sure that the Device server is running and that the local area network is available. If the problem persists, check the Device server log files for error messages that might help determine the problem.

BPCUI0019E No data is available for this selection.

Explanation

No data is available for this selection.

BPCUI0025E Probe job *job Id* failed.

Explanation

The probe job failed. The log will provide detailed information.

Action

BPCUI0029E Invalid parameter *param* passed.

Explanation

An invalid parameter was passed.

Action

BPCUI0030I This task was already executed.

Explanation

You are attempting to execute, schedule, or delete a task that was already executed. The task was possibly executed by another user.

Action

No action is required.

BPCUI0032E An unexpected response was received from the server.

Explanation

To find the cause of the issue, further investigation is required.

Action

To resolve the issue, try the following actions:

- Verify that the local area network is available.
- Check the status of the product servers on the Home > System Management page.
- Verify that the database repository is available.
- Verify that the related database service is active.
- Check for error messages in the log files for the servers.

If you need more information, go to the IBM Knowledge Center and check the Administering section. You can access the IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93>. If you cannot resolve the issue, contact IBM Support.

Related reference

-  [IBM Spectrum Control documentation](#)
-  [Getting support](#)

BPCUI0034E Invalid number of runs to keep for each schedule. The number should be between *param1* and *param2*.

Explanation

The entered number of runs to keep is not in the valid range. Please enter a valid number.

Action

BPCUI0035E Invalid number of days' worth of log-files to keep. The number should be between *param1* and *param2*.

Explanation

The entered number of days' worth of log-files to keep is not in the valid range. Please enter a valid number.

Action

BPCUI0036E The schedule id *scheduleID* associated with this job is no longer valid. It might have been deleted. Refresh the view and try again.

Explanation

The requested schedule does not exist in the database. It was either deleted in the meantime or the provided schedule id is not valid

Action

BPCUI0037E The replication server is not installed or is unavailable.

Explanation

The replication server cannot be reached. This condition might occur if the server is not installed or is unavailable.

Action

Verify that the replication server is installed and the local area network is available. If the server is installed, ensure that it is running and then try to add the resource again. If the server is not installed, run the product's installation program to install it and then try to add the resource again.

BPCUI0038E Invalid number of days to retain alerts. The number should be between *param1* and *param2*.

Explanation

The entered number day to retain alerts is not in the valid range. Please enter a valid number.

Action

BPCUI0039E A Storage Resource agent cannot be found.

Explanation

This problem might occur if a Storage Resource agent was not installed locally or was disabled.

Action

Verify that a Storage Resource agent is installed locally and is enabled. To determine if an agent is enabled on a server, go to the Servers page in the product, locate the server, and check the value in the Agent State column.

BPCUI0040E Parsing results from a call to the Data server failed with the following error message: *param1*.

Explanation

An issue was encountered while trying to parse results that were returned from a call to the Data server. This could be due to corrupted data sent from the data server, or to network problems, or other causes.

Action

Verify that the Data server is running properly and that the network is not experiencing difficulties.

BPCUI0042E Communication with the Data Server failed with the following error: *param1*

Explanation

Unable to contact the data server.

Action

Use the above error message to identify the cause of the problem. Verify that the data server, service, and port are all running properly and that database repository and the network are working properly.

BPCUI0043E Cannot connect to the Data server.

Explanation

The user interface cannot communicate with the Data server. This error might occur if the Data server is down, the Data server is in maintenance mode while updating agents or the local area network between the user interface and the Data Server is unavailable.

Action

Verify that the Data Server is up and running. Ensure that the Data server is not in maintenance mode and that the local area network is available. If the problem persists, check the Data server log files for error messages that might help determine the problem. See the product information center to view the locations of these log files.

If not already done, set the maximum level of tracing for the Data server to aid in resolution of the problem.

BPCUI0044E The entity was not found in the database.

Explanation

This error might occur if the entity was already deleted.

Action

BPCUI0045E Host name length exceeds the 255 character limit

Explanation

The host name must be less than or equal to 255 characters

Action

BPCUI0046E Report '*configurationId*' not found

Explanation

Post installation script may have failed.

Action

BPCUI0047E Parameter '*parameterName*' is not defined in report *configurationId*

Explanation

Parameter name could be incorrect or the parameter is not defined in report.

Action

BPCUI0048E No property is not defined for report *configurationId*'

Explanation

Parameter name could be incorrect or the parameter is not defined in report.

Action

BPCUI0049E No such property *propertyName* for report *configurationId*'

Explanation

Parameter name could be incorrect or the parameter is not defined in report.

Action

BPCUI0050E *variableName* can not be overridden

Explanation

Value for reportOutputFormat can be modified.

Action

BPCUI0051E *variableName* not valid report output format.

Explanation

PDF,XML,CSV,HTML,SINGLEXLS,SPREADSHEETML,XLWA are valid report output format.

Action

BPCUI0052E *variableName* not reachable

Explanation

TCR URL may not be reachable.

Action

BPCUI0053E Cannot authenticate with the provided user credentials.

Explanation

The user name or password that was entered for the device is not correct.

Action

Make sure that the user name and password are correct for the device that is being added. Reenter the user name or password and click Add again. If you are adding a Storwize V7000 Unified storage system, specify the IP address of the block component, not the filer component.

BPCUI0054E The host name or IP address {0} is not valid.

Explanation

The IP address or host name that was entered for the device is not valid.

Action

Make sure that the IP address and host name are valid for the device that is being added. Reenter the IP address or host name and click Add again.

BPCUI0055E Cannot connect to the storage system.

Explanation

This problem might be caused by the following conditions:

- For DS8000 storage systems, this problem might occur because the DS8000 ESSNI server is not available or is not allowing connections.
- For all storage systems, this problem might be caused by the values for connection protocols. IBM Spectrum Control uses different connection protocols to connect to resources. The default values of the connection protocols were changed for different releases of the products, which might cause connection failures.

Action

To resolve the problem, try the following actions:

- For DS8000 storage systems, verify that the ESSNI server is up and running. Use a tool such as ping to verify that the ESSNI server can be reached from the system where the Device server is installed.

If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)
- [Enabling and disabling legacy protocol \(SSLv3 and MD5 hash\)](#)

BPCUI0056E Cannot connect to the storage system or cluster.

Explanation

The device that is being added might not be up and running. Also, the IP address, host name, or user credentials that was entered for the storage system or cluster might not be valid.

Action

Verify that the device is up and running. Also, make sure that the IP address, host name, and user credentials are correct for the device that is being added. Reenter the values and click Add again.

BPCUI0058I No supported resources were discovered on the data source *data_Source_Address*.

Explanation

The resources that were discovered on the specified data source are not supported.

Action

For a list of data sources that you can use with the product, go to the support matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

BPCUI0060I File *param* was successfully uploaded to the Data Server.

Explanation

The file the user selected was successfully uploaded to the Sata Server.

Action

The file should be placed on the Data Server and ready to be used by the deployment job.

BPCUI0061E Upload file type *param* is not supported.

Explanation

The upload file type is not supported.

Action

Check the upload file type and provide a correct file type.

BPCUI0062E The requested action failed with the following error message: *error message*

Explanation

The requested action did not complete, and the specified error message was returned.

Action

BPCUI0063E Cannot find jobs for *scheduleId param* and *deviceId param*. No logs are displayed.

Explanation

Jobs related to the schedule and device are not available.

Action

Verify that the schedule was run for the device.

BPCUI0064E A log file cannot be displayed for the job.

Explanation

The selected job does not have an associated job log file.

Action

Verify that the database repository is up and the related database service is active. Ensure that all required product servers are running and that the local area network is available. To check the status of the servers, go to the Home > System Management page.

You can also check the log files of the servers for error messages that might help determine the problem.

For information about how to determine if the database repository is active, if the product servers are running, and the location of log files, go to the IBM Knowledge Center and check the Administering section. You can access the IBM Knowledge Center for the product at <http://www.ibm.com/support/knowledgecenter/SS5R93>.

BPCUI0065E The job log file cannot be accessed. The log file may have been manually removed or may have been deleted because it was older than *retain_days* days or it exceeded the maximum number of *no_of_lofs* runs.

Explanation

The job log file cannot be accessed. The log file may have been manually removed or may have been deleted because it was older than the specified number of days or it exceeded the maximum number of runs.

Action

Verify that the log file exists in the specified location.

BPCUI0067E The schedule for collecting status and asset data cannot be created.

Explanation

If you can't schedule a probe to collect data about a resource, it might mean that the product servers are down or the local area network is not available.

Action

Ensure that all required product servers are running and that you have a network connection to the system on which they are located. To check the status of the servers, go to the Home > System Management page.

BPCUI0068E A proposed schedule for collecting status and asset data cannot be created.

Explanation

No further explanation required.

Action

Ensure that all required product servers are running and that you have a network connection to the system on which they are located. To check the status of the servers, go to the Home > System Management page.

BPCUI0069E The proposed schedule for collecting status and asset data cannot be deleted.

Explanation

No further explanation required.

Action

Ensure that all required product servers are running and that you have a network connection to the system on which they are located. To check the status of the servers, go to the Home > System Management page.

BPCUI0071E The task *task_name* could not be completed.

Explanation

An error prevented the task from completing. For further details, consult the system log file.

Action

Check the trace.log file in the logs directory for entries mentioning the failed task. If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCUI0072E Cannot connect to the Device server. Verify that the database service and Device server are running, and that the Device server is accessible.

Explanation

The Device server is unavailable. This error might occur if the Device server is down or the local area network is unavailable.

Action

Verify that the database service and Device server are up and running. Ensure that you have a network connection to the server on which the Device server is located. Try the action again.

BPCUI0073E Can't make a connection to the *storage_resource* storage resource.

Explanation

The connection to the resource with the host name or IP address that was specified was unsuccessful.

Action

Complete the following tasks:

- Check that the correct host name or IP address of the storage resource was entered.
- Check whether you have a network connection to the storage resource that was specified.
- Check whether the server that hosts the product is up and running.

BPCUI0074E The wizard could not set an attribute for the storage resource.

Explanation

The wizard could not set a user-defined attribute for the storage resource because an error occurred on the Device server. Examples of user-defined attributes include the location and display name of the storage resource.

Action

Ensure that the Device server is running and that the local area network is available. If the problem persists, check the Device server log files for error messages that might help determine the problem. See the product information center to view the locations of the log file.

BPCUI0075E The certificate wasn't saved on the server.

Explanation

Certificates are saved to a repository called a truststore. The certificate for this device could not be saved to the truststore.

Action

Ensure that all required product servers are running and that you have a network connection to the system on which they are located. To check the status of the servers, go to the Home > System Management page.

BPCUI0076W The initial job to collect status and asset data did not start.

Explanation

The wizard cannot start a job for collecting status and asset data from the storage resource.

Action

Ensure that the Device server is running and that the local area network is available. If the first scheduled job to collect status and asset data also does not start, check the Device server log files and Data server log files for error messages that might help determine the problem. See the product information center to view the locations of these log files.

BPCUI0077E A failure occurred loading the certificate.

Explanation

Loading of a certificate to the truststore failed for an unknown reason.

Action

Contact IBM

Related reference

- [Getting support](#)

BPCUI0078I The certificate was loaded successfully.

Explanation

The user has successfully loaded a certificate into the server's truststore.

Action

No action is required.

BPCUI0079E The SSL certificate is not in the expected format.

Explanation

Loading of a certificate to the truststore failed because the certificate was not in the x.509 format.

Action

Ensure that the certificate being loaded is in the X.509 format.

BPCUI0084W The wizard could not retrieve the default interval information for performance monitoring.

Explanation

The wizard could not retrieve the default interval information for the device from the Device server. The interval value represents the number of minutes over which performance data is averaged.

Action

Ensure that the Device server is running and that the local area network is available. Check the Device server log files for error messages that might help determine the problem. See the product information center to view the locations of these log files.

BPCUI0085E The user name or password for the *hypervisor or vCenter* hypervisor or vCenter Server is invalid.

Explanation

The user name or password is not valid for the specified hypervisor or vCenter Server.

Action

To resolve the problem, try the following actions:

- Ensure that you specified the correct host name, user name, and password for the hypervisor or vCenter Server.
- Verify that the local area network is available and that you can connect to the hypervisor or vCenter Server.
- Ensure that the Device server is running.
- Check the Device server log files for error messages that might help to determine the problem. For the location of these log files, search the IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93/>.

BPCUI0086E The SSL certificate is invalid for the *hypervisor or vCenter hypervisor or vCenter Server*, or the firewall is blocking access to it.

Explanation

The product cannot communicate with the hypervisor or vCenter Server.

Action

To resolve the problem, try the following actions:

- Ensure that the SSL certificate file is valid for the hypervisor or vCenter Server.
- Verify that the local area network is available and that you can connect to the hypervisor or vCenter Server.
- Ensure that any required firewall authorization was granted.
- Ensure that the Device server is running.
- Check the Device server log files for error messages that might help to determine the problem.

If you need more information, go to the IBM Knowledge Center and check the "Administering" section. You can access the IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93/welcome>.

BPCUI0087E The version of the *hypervisor or vCenter hypervisor or vCenter Server* is not supported.

Explanation

You can manage only the supported versions of hypervisors and vCenter Servers.

Action

Check the product support site at www.ibm.com/support/docview.wss?uid=swg27039840 for a list of hypervisor and vCenter Server versions that are supported.

BPCUI0088E The host name, protocol, or port for the *hypervisor or vCenter hypervisor or vCenter Server* is invalid, or the hypervisor or vCenter Server is unreachable.

Explanation

The specified host name, protocol, or port number is invalid and cannot be used to communicate with the hypervisor or vCenter Server, or the hypervisor or vCenter Server is unreachable.

Action

To resolve the problem, try the following actions:

- Ensure that you specified the correct host name, protocol, and port number for the hypervisor or vCenter Server.
- Verify that the local area network is available and that you can connect to the hypervisor or vCenter Server.
- Ensure that the Device server is running.
- Check the Device server log files for error messages that might help to determine the problem. For the location of these log files, search the IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93/>.

BPCUI0089W Cannot retrieve a valid set of data collection intervals for performance monitoring.

Explanation

Monitored devices have a specific set of intervals that determine how often their performance data can be collected. For this device, the interval information could not be retrieved from the Device server.

Action

Ensure that the Device server is running and that the local area network is available. Check the Device server log files for error messages that might help determine the problem. See the product information center to view the locations of these log files.

BPCUI0090I All alerts were removed.

Explanation

This message is for informational purposes only

Action

No further action is required

BPCUI0091W *error_count* of *total_count* alerts were not removed.

Explanation

>Not all the acknowledged alerts were removed in the alert log..

Action

>Ensure that the local area network is available. Verify that the Device server and Data server are running, and that the database is operational. Check the product log files for error messages that might help determine the problem. See the product information center to view the locations of these log files.

BPCUI0093I No data path is available for *deviceNameVariable*.

Explanation

The selected device either does not have a data path or the devices for the data path are not known to the system.

Action

There is no action if the device does not have a data path. If the device is part of a data path, ensure the other devices are added to the system and probed.

BPCUI0094E Authorization failed due to an internal error.

Explanation

The authorization infrastructure was not initialized successfully.

Action

Contact IBM support.

Related reference

-  [Getting support](#)

BPCUI0097E Authorization failed due to an invalid request context.

Explanation

The authorization infrastructure was not initialized successfully.

Action

Contact IBM support.

Related reference

-  [Getting support](#)

BPCUI0098E The current user is not authorized to perform the requested function.

Explanation

The role assigned to the current user does not have sufficient privileges to perform the requested function.

Action

Request additional privileges from the storage administrator.

BPCUI0099E The storage resource is not available.

Explanation

The storage resource is not being monitored and cannot be displayed. This problem might occur if a bookmark in the web browser was used to access the page for the storage resource, but that resource is no longer being monitored.

This problem might also occur if the product servers are down or the local area network is not available.

Action

To resolve the problem, try the following actions:

- Remove the bookmark for the storage resource from the web browser.
- Monitor the resource again. To add the storage resource for monitoring, go to the Storage Systems page in the GUI and select Add Storage System. See the online help for information about how to add a storage resource for monitoring.

If the storage resource is being monitored and you receive this error message, try the following actions:

- Ensure that all required product servers and services are running and that you have a network connection to the system on which they are located. To check the status of the servers, go to the Home > System Management page.
- Check for error messages in the log files for the servers.

If you need more information, go to the IBM Knowledge Center and check the Administering section. You can access the IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93>. If the problem continues, contact IBM Software Support.

Related reference

-  [Getting support](#)

BPCUI0100I *success_count* alerts were marked as acknowledged.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BPCUI0101I The alert was marked as acknowledged.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BPCUI0102E None of the alerts were marked as acknowledged.

Explanation

None of the alerts in the alert log were marked as acknowledged.

Action

Verify that there are unacknowledged alerts in the list. Ensure that the local area network is available, and that the Device server and Data server are running, and that the database is operational. Check the product log files for error messages that might help determine the problem. See the product information center to view the locations of these log files.

BPCUI0104I *success_count* alerts were marked as unacknowledged.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BPCUI0105I The alert was marked as unacknowledged.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BPCUI0108I All informational alerts were marked as acknowledged.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BPCUI0110W Some informational alerts were not marked as acknowledged.

Explanation

Not all of the informational alerts were marked as acknowledged in the alert log.

Action

Ensure that the local area network is available. Verify that the Device server and Data server are running, and that the database is operational. Check the product log files for error messages that might help determine the problem. See the product information center to view the locations of these log files.

BPCUI0111I All alerts were marked as acknowledged.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BPCUI0112I *success_count* alerts were removed.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BPCUI0113I The alert was removed.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BPCUI0114I All acknowledged alerts were removed.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BPCUI0116W Some acknowledged alerts were not removed.

Explanation

Not all of the acknowledged alerts were removed from the alert log.

Action

Ensure that the local area network is available. Verify that the Device server and Data server are running, and that the database is operational. Check the product log files for error messages that might help determine the problem. See the product information center to view the locations of these log files.

BPCUI0120W Some acknowledged alerts were not marked as unacknowledged.

Explanation

Not all the acknowledged alerts were changed to unacknowledged in the alert log.

Action

Ensure that the local area network is available. Verify that the Device server and Data server are running, and that the database is operational. Check the product log files for error messages that might help determine the problem. See the product information center to view the locations of these log files.

BPCUI0121E Unable to communicate with the product server. Make sure that the server is running properly.

Explanation

The most recent request that was sent to the Web server did not complete. This problem might occur if communication with the server cannot be established, or if the local area network is unavailable.

Action

To resolve this problem, try the following actions:

- Go to the Home-System Management page in the GUI and ensure that the Web server is active.
- Verify that the local area network is available.
- Verify that you have a network connection to the computer on which the server is located.
- Check the log files for error messages that might help determine the problem. If you need more information, go to IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93> and check the Administering section.

If the problem continues, contact IBM Software Support.

Related reference

-  [Getting support](#)

BPCUI0122E No job log file was created for this job run.

Explanation

The server did not create a job log for the job run.

Action

Ensure that all required product servers and services are running and that you have a network connection to the system on which they are located. To check the status of the servers, go to the Home > System Management page.

BPCUI0123E The action cannot be completed.

Explanation

An error occurred when the server tried to process a request. This message might occur when a request is not handled correctly by the internal framework.

Action

Ensure that all required product servers and services are running and that you have a network connection to the system on which they are located. To check the status of the servers, go to the Home > System Management page.

Check the log files of the servers and web-based GUI for error messages that might determine the problem. See the product documentation for the location of the log files.

BPCUI0124E An unexpected error occurred during the execution of the action.

Explanation

While executing the action, an unexpected error occurred.

Action

Contact IBM Support.

Related reference

- [Getting support](#)

BPCUI0125E The alert is not available.

Explanation

The alert cannot be displayed on the page. This problem might occur if a bookmark or URL was used to access the alert, but that alert has been removed.

Action

Remove the bookmark for the storage resource from the web browser.

BPCUI0126E The status of the Performance Monitors could not be retrieved.

Explanation

The status of the Performance Monitors could not be retrieved from the Device Server.

Action

Verify that the Device Server is running properly and that the network is not experiencing difficulties.

BPCUI0127E The currently installed version of the product does not have the required product license for the function that you requested.

Explanation

You need a different product license to perform the function that you requested.

Action

For information about how to upgrade to the product license that you need, go to IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93>.

BPCUI0128E An undefined capacity chart metric was requested.

Explanation

While viewing a capacity bar chart, an unknown metric was requested. This error should never be seen by the user.

Action

Customer should not see this error. Contact IBM Support.

Related reference

- [Getting support](#)

BPCUI0129I Alerts that were migrated from a previous version of the product are not shown on this page.

Explanation

Alerts are generated when certain conditions are detected on monitored resources. The version of the product that detected an alert determines where information about the alert is displayed. Alerts that were generated in recent versions of the product are displayed on the Alerts page and also on the detail pages of their associated resource. Alerts that were generated in older versions of the product, also called migrated alerts, are only displayed on the Alerts page.

Action

If you migrated alerts from a previous version, go the Alerts page to view the full list of all historical alerts for the resource that is being displayed. To access the Alerts page, click the link in the message or select Alerts from the menu.

This warning message is shown only if one or more migrated alerts still exist in the database. When all migrated alerts are either manually deleted or become old enough to be deleted automatically, this warning is no longer displayed. Alternatively, you can suppress the warning message by clicking Don't show this again in the warning message window.

BPCUI0130E The alerts cannot be acknowledged because they were deleted.

Explanation

None of the alerts were marked as acknowledged in the alert log because these no longer exist in the database.

Action

No further action is required.

BPCUI0131E The alerts cannot be unacknowledged because they were deleted.

Explanation

None of the alerts were marked as unacknowledged in the alert log because these no longer exist in the database.

Action

No further action is required.

BPCUI0132W *success_count* alerts were marked as acknowledged. *unsuccess_count* alerts cannot be marked as acknowledged because they were deleted.

Explanation

Not all of the alerts were marked as acknowledged in the alert log because these no longer exist in the database.

Action

No further action is required.

BPCUI0133W *success_count* alerts were marked as unacknowledged. *unsuccess_count* alerts cannot be marked as unacknowledged because

they were deleted.

Explanation

Not all of the alerts were marked as unacknowledged in the alert log because these no longer exist in the database.

Action

No further action is required.

BPCUI0134E The alert cannot be acknowledged because it was deleted.

Explanation

The alert was not marked as acknowledged in the alert log because this no longer exist in the database.

Action

No further action is required.

BPCUI0135E The alert cannot be unacknowledged because it was deleted.

Explanation

The alert was not marked as unacknowledged in the alert log because this no longer exist in the database.

Action

No further action is required.

BPCUI0136E The device was not removed because the action is not supported for devices of type *devType*.

Explanation

An internal operating error occurred. Check the logs for an indication of an error or exception and contact IBM customer support.

Action

Check the logs for an indication of an error or exception and contact IBM customer support.

Related reference

- [Getting support](#)

BPCUI0137E Input text provided has invalid character(s): *characters*. Input text: *text*

Explanation

Action

Please do not use the characters which are not allowed in the input.

BPCUI0141E Host name or IP address *hostname* specified on line *line* of file *file* is not valid.

Explanation

The specified host name or IP address in the file is not valid.

Action

Enter a valid host name or IP address. Host names must be less than 255 characters in length, contain letters 'a' through 'z' or digits '0' through '9'. Hyphens are allowed as long as it is not the leading or trailing character in the host name.

BPCUI0143E Host port WWPN *wwpn* specified on line *line* of file *file* is not valid.

Explanation

The specified WWPN in the file is not valid.

Action

Enter a valid host port WWPN. WWPNs must be 16 hexadecimal characters in length with no other characters.

BPCUI0144E Duplicate server *name* specified on lines *line1* and *line2* of file *file*.

Explanation

The file contains duplicate entries for the same server.

Action

Ensure the file contains one entry per server to create.

BPCUI0145E Could not parse file *file*.

Explanation

The attempt to parse the file failed.

Action

Verify the correct file was specified or the file is in the correct format.

BPCUI0146E Could not parse file *file*. Invalid entry on line *line*.

Explanation

The attempt to parse the file failed.

Action

Verify the correct file was specified or the file is in the correct format.

BPCUI0148I Successfully deleted server *server_name*.

Explanation

The server was successfully deleted.

Action

No action is required.

BPCUI0149I Successfully modified ports of server *server_name*.

Explanation

The ports of the server were successfully modified.

Action

No action is required..

BPCUI0150I The server was created.

Explanation

The server was created.

Action

The server was created.

BPCUI0151E The host name or IP address is associated with another resource.

Explanation

The same host name or IP address cannot be used to connect to different resources.

Action

Enter a host name or IP address that is not associated with a resource that was already added for monitoring.

BPCUI0152I The data source *data_Source_Address* was successfully added as a data source for monitoring. The following new resources were detected:

Explanation

You cannot configure resources that are already being monitored. Only resources that are newly discovered or are not included in a data collection schedule can be added for monitoring.

Action

No action is required.

BPCUI0155W You cannot provision volumes because there is no Fibre Channel host port information for at least one server.

Explanation

To provision a volume to a server by using the Provision Storage wizard, the server must have at least one Fibre Channel host port. If you are provisioning to a single server, the server does not have a Fibre Channel host port. If you are provisioning to multiple servers, at least one of the servers does not have a Fibre Channel host port.

Action

To request volumes by using the Provision Storage wizard, ensure that all the selected servers have Fibre Channel host ports.

BPCUI0156W You cannot provision volumes to servers that use different operating systems.

Explanation

To provision a volume to multiple servers by using the Provision Storage wizard, the servers must all be running the same operating system type.

Action

To provision volumes to multiple servers, make sure that you select servers that all run the same operating system type. The OS Type column on the Servers page displays the operation system type.

BPCUI0157W You cannot provision volumes to servers and virtual machines at the same time. To provision volumes, ensure that you select either only servers or only virtual machines.

Explanation

By using the Provision Storage wizard, you can provision volumes to multiple servers or to multiple virtual machines. You cannot, however, provision to a combination of servers and virtual machines.

Action

When you open the Provision Storage wizard from the Servers page, make sure that you select only servers or only virtual machines. The Virtual Machines column on the Servers page indicates whether the server is a virtual machine.

BPCUI0158I Volumes are assigned to the hypervisors that host virtual machines. Volumes are not assigned directly to virtual machines.

Explanation

When you provision volumes for one or more virtual machines that do not have NPIV ports, the volumes are assigned to the hypervisors that host the virtual machines.

Action

To assign volumes to virtual machines, wait until the provisioning operation completes. Then, use the hypervisor manager to assign the hypervisor disks to the virtual machines.

BPCUI0159W You cannot provision volumes because at least one of the hypervisors that host the virtual machines is not being monitored. Ensure that all the hypervisors that are hosting the virtual machines that were selected for provisioning were probed.

Explanation

When you provision volumes for one or more virtual machines that do not have NPIV ports, the volumes are assigned to the hypervisors that host the virtual machines. If you are requesting storage for a single virtual machine, that hypervisor that hosts the virtual machine is not being monitored. If you are requesting storage for multiple virtual machines, one or more of the hypervisors that are hosting the virtual machines are not being monitored.

Action

When you request volumes for virtual machines that do not have NPIV ports, make sure all of the hypervisors that host the virtual machines were probed.

BPCUI0160E Duplicate port WWPN *wwpn* specified on lines *line1* and *line2* of file *file*.

Explanation

The file contains duplicate entries for the same server port.

Action

Ensure the file contains ports that are not assigned to more than one server.

BPCUI0162W File *file* does not contain any servers to create.

Explanation

The specified file does not contain any servers to create.

Action

Specify a file which includes entries for servers to create.

BPCUI0166W Optimization cannot be done in place to the subsystem since storage subsystem *param1* and/or its pools belong to more than one capacity pool. Following are capacity pools the subsystem is associated with: *param2*

Explanation

Optimization cannot be done in place to the subsystem since storage system and/or storage pools from the system belong to more than one capacity pools.

Action

Specify a target capacity pool for optimization or if you want to use in place optimization, refer to the documentation for proper configuration.

BPCUI0167W Optimization cannot be done in place to the subsystem since storage subsystem *param1* and/or its pools are not part of any capacity pool.

Explanation

Optimization cannot be done in place to the subsystem since storage system and/or storage pools from the system do not belong to any capacity pool.

Action

Specify a target capacity pool for optimization or if you want to use in place optimization, refer to the documentation for proper configuration.

BPCUI0168W Optimization cannot be done in place to the server *param1* since storage subsystems or storage pools associated with luns assigned to the server belong to more than one capacity pool. Following are associated capacity pools: *param2*

Explanation

Optimization cannot be done in place to the server since storage systems and/or storage pools associated with luns assigned to the server belong to more than one capacity pools.

Action

Specify a target capacity pool for optimization or if you want to use in place optimization, refer to the documentation for proper configuration.

BPCUI0169W Optimization cannot be done in place to the server *param1* since storage subsystems or storage pools associated with luns assigned to the server are not part of any capacity pool.

Explanation

Optimization cannot be done in place to the server since storage system and/or storage pools associated with luns assigned to the server do not belong to any capacity pool.

Action

Specify a target capacity pool for optimization or if you want to use in place optimization, refer to the documentation for proper configuration.

BPCUI0170W Optimization cannot be done in place to the storage entity *param1* since storage subsystems or storage pools associated with it belong to more than one capacity pool. Following are associated capacity pools: *param2*

Explanation

Optimization cannot be done in place to the entity since storage systems and/or storage pools associated belong to more than one capacity pools.

Action

Specify a target capacity pool for optimization or if you want to use in place optimization, refer to the documentation for proper configuration.

BPCUI0171W Optimization cannot be done in place to the storage entity *param1* since storage subsystems or storage pools associated with it are not part of any capacity pool.

Explanation

Optimization cannot be done in place to the entity since storage system and/or storage pools associated with it do not belong to any capacity pool.

Action

Specify a target capacity pool for optimization or if you want to use in place optimization, refer to the documentation for proper configuration.

BPCUI0172E The operation timed out while waiting for a response from the server.

Explanation

The server did not respond to a request from the GUI. This message is displayed if the server does not respond within the length of time defined for the timeout value.

Action

Ensure that all required servers are running, and that the local area network is available. Verify that you have a network connection to the server that the database repository is located on. Verify that the database repository is up and the related database service is active.

BPCUI0173E File *file* does not exist or is empty.

Explanation

The file does not exist in the specified location or is empty.

Action

Either create the file or select a new one with servers to create.

BPCUI0174E The device does not support the credential mechanism used.

Explanation

The device does not support the credential mechanism used.

Action

Choose another credential mechanism.

BPCUI0175E A required parameter is missing.

Explanation

The highlighted field is missing a required value. As a result, a required parameter in the server request was missing or empty.

Action

Enter a valid value in the highlighted field.

BPCUI0176E The highlighted field contains an invalid value.

Explanation

The highlighted field contains an invalid value. As a result, a parameter in the server request contained an invalid value.

Action

Enter a valid value in the highlighted field.

BPCUI0177E The highlighted field contains a value that is outside of the allowed range. The value must be between *minVal* and *maxVal*.

Explanation

The highlighted field contains a value that is outside of the allowed range. As a result, a parameter in the server request is outside of the allowed range.

Action

In the highlighted field, enter a value that is within the allowed range.

BPCUI0178E A service class with the same name and type already exists.

Explanation

The specified service class name is not unique. A service class with the same name and type already exists in the database and cannot be duplicated.

Action

Enter a different name for the service class you are creating.

BPCUI0179I The service class was created.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCUI0180I Based on the known configuration of storage system host connections, fabric zone aliases, and HBA ports, additional ports may have been added to the selection below.

Explanation

Based on the input host name or WWPN, additional ports were found in the storage system host connection configuration, fabric zone alias configuration, or HBA port configuration.

Action

Remove unwanted WWPNs by clicking the remove icon. Add additional WWPNs by typing them in then clicking the Add button.

BPCUI0181I You selected to add a *expectedDevice* resource, but a *foundDevice* resource was detected and will be added.

Explanation

One or more resources that are managed by the data source are of a different type than expected. However, the resource will still be added for monitoring if you enter its information. For example, if you select to add a SAN Volume Controller, but the discovered resource is a Storwize V7000, this message is displayed. In this case, the Storwize V7000 storage system will be added.

Action

To add the resource that was discovered, continue to enter its information. To add the intended resource, click Add Storage System again and enter the correct IP address for the resource.

BPCUI0182I The data source *data_Source_Address* was added as a data source for monitoring. No new resources were detected.

Explanation

Only resources that are newly discovered or are not included in a data collection schedule can be added for monitoring.

Action

Enter the IP address or host for a different data source or resource to continue.

If you want to modify resources that are already monitored, go to the list page for the resource. For example, to modify a storage system, in the navigation pane select Storage Resources > Storage Systems. Then, right-click the resource and select View Properties. In the properties notebook, modify the values for the resource.

BPCUI0183E The text in the highlighted field exceeds the *maxLength* character limit.

Explanation

The text in the highlighted field is too long.

Action

Enter text in the highlighted field that does not exceed the character limit.

BPCUI0185W Unable to lookup the IP Address for Host Name *hostName*. Enter the IP Address manually.

Explanation

Unable to lookup the IP Address for the given host name.

Action

Manually enter the IP Address.

BPCUI0189I Configuration of SRA deployment and probe schedules were done successfully.

Explanation

Configuration of deployment probe schedules were done successfully.

Action

None.

BPCUI0190W Configuration of SRA finished with some warnings or errors. Check the detail messages.

Explanation

Configuration user trying to make finish with warnings or errors..

Action

Check the accompanying message to take further actions.

BPCUI0191E An internal error occurred while testing conneciton to *param1*.

Explanation

This type of error is usually caused by unexpected run time condition like when Data/Device Server is down or communication to the servrs are interrupted unexpectedly.

Action

Please check the log file for the errors. Correct them and retry the operation.

BPCUI0192E The supplied service class type is invalid.

Explanation

An Invalid service class type was passed to the server.

Action

Contact IBM Software support.

Related reference

- [Getting support](#)

BPCUI0193E The specified SMI-S provider was not found. Make sure that the protocol, SMI-S provider host name or IP address, and port are specified correctly and that the SMI-S provider is properly configured at that location.

Explanation

The SMI-S provider specified was not found.

BPCUI0194E An unknown error has occurred. Please review all values entered.

Explanation

An unknown error has occurred.

BPCUI0195E The Interop Namespace is not correct. Please correct this entry.

Explanation

The Interop Namespace is not correct.

BPCUI0196E A timeout occurred while processing the request. Please retry request.

Explanation

The request could not be processed in the time allowed.

BPCUI0197E A connection was not established. Make sure that the protocol, SMI-S provider host name or IP address, and port are specified correctly.

Explanation

The attempt to establish a connection failed.

Related reference

- [Resolving communication issues with Brocade Network Advisor](#)

BPCUI0198E The authentication to the SMI-S provider failed.

Explanation

The credentials that you supplied for the connection are incorrect.

BPCUI0199E An `SSLHandshakeException` or `SSLProtocolException` has occurred. This exception might be due to an invalid SLP registration, e.g. 'http' instead of 'https'.

Explanation

An `SSLHandshakeException` or `SSLProtocolException` has occurred.

BPCUI0201E There is a pending delete in process for this SMI-S provider.

Explanation

The specified SMI-S provider is currently being deleted.

BPCUI0202I Success

Explanation

The operation was successful.

BPCUI0203E The selected resources were not removed.

Explanation

None of the resources were removed.

Action

Review the listed of detailed messages for further information.

BPCUI0204W *successfulDeletes* of *attemptedDeletes* of the selected resources were removed.

Explanation

Removing the resources were partially successful, but some warnings occurred.

Action

Review the listed of detailed messages for further information.

BPCUI0205W *successfulDeletes* selected resources were removed, however warnings did occur.

Explanation

The resources were removed however some warnings did occur.

Action

Review the listed of detailed messages for further information.

BPCUI0209E A database operation cannot be completed.

Explanation

An SQL error was encountered when attempting to complete an action on the database repository.

Action

Verify that that the database repository is running properly and the related database service is active. Ensure that all required servers are running and that the local area network is available. Verify that you have a network connection to the server on which the database repository is located. To check the status of the servers, go to the Home > System Management page. Check the log files of the servers for error messages that might help determine the problem. See the product information center for the location of these log files.

BPCUI0210I Device *param1* supports performance monitoring.

Explanation

Performance monitoring is supported on this device.

Action

BPCUI0211E No performance data is available for a resource.

Explanation

A resource is not being managed by the storage management service, so no performance data is available for that resource.

Action

To monitor the storage resource, go to the Storage Systems page in the GUI and select Add Storage System. See the online help for information about how to add a storage resource for monitoring.

BPCUI0212E There is no Secure Shell running at this host/IP.

Explanation

SSH is not running on this host/IP.

Action

Enable SSH on this host/IP and make sure it is running

BPCUI0213E Unsupported Secure Shell protocol was used.

Explanation

Action

BPCUI0214E Invalid public key location for subsystem *param1*.

Explanation

Action

BPCUI0215E Invalid public key format for subsystem *param1*.

Explanation

Action

BPCUI0216E Passphrase was incorrect for subsystem *param1*.

Explanation

Action

BPCUI0217E Unable to transfer the key(s) to the server *param1*.

Explanation

Action

BPCUI0218E The specified private key file format is not supported. Please convert it to Open SSH (.pem) key file format for subsystem *param1*.

Explanation

Action

BPCUI0219E The specified key file or key file name is already linked to another user.

Explanation

Action

BPCUI0220E The IP address that was entered was the address of the management console for the storage system. You must enter the valid IP address of the block component of the storage system.

Explanation

Action

Enter the IP address of the block component. Try to add the storage system again.

BPCUI0221E The IP address you entered is the address of another device's management console.

Explanation

Action

BPCUI0222E The IP address you entered points to a device of another type.

Explanation

This error might occur in situations such as the following:

- You are trying to add a storage system of a certain type, but you enter the IP address for a different type of storage system
- You are trying to add a storage system, but you enter the IP address of a switch
- You are trying to add a storage system, for example IBM Spectrum Scale, but the storage system software is not installed on the device

Action

To resolve the problem, try the following actions:

- Ensure that the host name or IP address is valid for the device type that you are adding
- Ensure that the device is operational
- Try to add the device again

BPCUI0223E Passphrase is required. Specify one for subsystem *param1*.

Explanation

The passphrase for the truststore is missing.

Action

Provide the missing passphrase

BPCUI0224E Cannot connect to a resource because of an SSL certificate error. Troubleshooting information: <http://www.ibm.com/support/docview.wss?uid=swg21976237>

Explanation

This communication problem might be caused by an error with the SSL certificate on the resource.

Action

To learn about how to troubleshoot the problem, go to <http://www.ibm.com/support/docview.wss?uid=swg21976237>.

BPCUI0225I The agent log files for *server_Name* have been collected and copied to *log_Location*.

Explanation

The server successfully collected the Storage Resource agent log files and saved them to the specified location.

BPCUI0226I Discovery of *data_source* is taking longer than expected. Click Close to run the discovery in the background.

Explanation

The process for discovering the resources that are managed by the data source is taking longer than expected. This delay might occur because the network connection is slow, the data source is running slowly, or the data source manages many resources.

Action

To close the window but continue the discovery process in the background, click Close. Try to add the resource again when the discovery is complete.

BPCUI0227E Thin provisioning must be enabled when compression is enabled.

Explanation

The service class cannot be created because compression is enabled but thin provisioning is disabled. It is not possible to create a compressed volume that is not also thin provisioned.

Action

Enable thin provisioning when compression is enabled.

BPCUI0229I 1 resource was added to *name*.

Explanation

The indicated number of resources are now members of the Capacity Pool.

Action

None.

BPCUI0231I *count* resources were added to *name*.

Explanation

The indicated number of resources are now members of the Capacity Pool.

Action

None.

BPCUI0233E The specified host name is already associated with an existing server.

Explanation

You cannot add a server with same host name as a server that is already being monitored.

Action

Enter a host name for a server that is not already being monitored.

BPCUI0234E The specified IP address is already associated with an existing server.

Explanation

You cannot add a server with same IP address as a server that is already being monitored.

Action

Enter an IP address for a server that is not already being monitored.

BPCUI0235E The specified host name and IP address are already associated with an existing server.

Explanation

You cannot add a server with same IP address and host name as a server that is already being monitored.

Action

Enter a host name and IP address for a server that is not already being monitored.

BPCUI0236E The disabling of the agents failed.

Explanation

None of the selected agents were disabled due to errors.

Action

Review the individual detailed error messages.

BPCUI0237E Errors occurred when attempting to disable some of the agents.

Explanation

Disabling some of the agents was successful, however errors did occur on others.

Action

Review the individual detailed error messages.

BPCUI0238W Warnings occurred when attempting to disable *warningCount* of the agents.

Explanation

The agents were successfully disabled, but warnings did occur

Action

Review the individual detailed warning messages.

BPCUI0239I *attemptedCount* of the *selectedCount* selected agents were disabled.

Explanation

All agents that could be disabled did so successfully.

Action

No action.

BPCUI0240E The agents were not enabled.

Explanation

None of the selected agents were enabled due to errors.

Action

Review the individual detailed error messages.

BPCUI0241E Errors occurred when attempting to enable some of the agents.

Explanation

Enabling some of the agents was successful, however errors did occur on others.

Action

Review the individual detailed error messages.

BPCUI0242W Warnings occurred when attempting to enable *warningCount* of the agents.

Explanation

The agents were successfully enabled, but warnings did occur.

Action

Review the individual detailed warning messages.

BPCUI0243I *attemptedCount* of the *selectedCount* selected agents were enabled.

Explanation

All agents that could be enabled did so successfully.

Action

No action.

BPCUI0244I The credentials of an agent were updated.

Explanation

The agent credentials were successfully updated

Action

No action required.

BPCUI0245I The credentials of *updateCount* agents were updated.

Explanation

The agent credentials were successfully updated

Action

No action required.

BPCUI0246E Cannot authenticate to the file module with the provided user credentials.

Explanation

The user name or password that was entered for the file module is not correct.

Action

Make sure that the user name and password are correct for the file module that is being added. Reenter the user name or password and click Add again.

BPCUI0247E Unknown file module key user.

Explanation

The key user that was entered for the file module was not found.

Action

Make sure that the key user is correct for the file module that is being added. Reenter the key user and click Add again.

BPCUI0248E The SSH key could not be loaded for the following reason: *IOException message*

Explanation

Upload of the file failed. Please retry the upload.

BPCUI0249E Passphrase is incorrect.

Explanation

Action

BPCUI0250E Passphrase is required.

Explanation

The passphrase for the truststore is missing.

Action

Provide the missing passphrase

BPCUI0251E Cannot connect to the storage system or cluster.

Explanation

The device that is being added might not be up and running. Also, the IP address, host name, or user credentials that was entered for the storage system or cluster might not be valid.

Action

Verify that the device is up and running. Also, make sure that the IP address, host name, and user credentials are correct for the device that is being added. Reenter the values and click Add again.

BPCUI0252E The host name or IP address {0} is not valid.

Explanation

The IP address or host name that was entered for the device is not valid.

Action

Make sure that the IP address and host name are valid for the device that is being added. Reenter the IP address or host name and click Add.

BPCUI0253E Cannot connect to the data source for the resource with the address *ip_address*.

Explanation

The data source for the storage resource cannot be contacted.

BPCUI0254E Invalid private key location.

Explanation

BPCUI0255W The following resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?

Explanation

You are attempting to add one or more resources to a capacity pool, but the resources are already assigned to another capacity pool. If you continue, the resources will be reassigned to the capacity pool you specified.

Action

BPCUI0256W The following resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?

Explanation

When a storage system is added to a capacity pool, all of its internal resources (storage pools, NSDs, and file systems) are indirectly assigned to that capacity pool. You are attempting to add one or more resources to a capacity pool, but the resources are already indirectly assigned to another capacity pool at the storage system level. If you continue, the resources will be reassigned to the capacity pool you specified. In addition, the storage system will no longer be assigned to its capacity pool. The storage system resources that are not being reassigned will instead be directly assigned to the storage system's original capacity pool.

BPCUI0257W The following resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?

Explanation

You are attempting to add one or more resources to a capacity pool, but the resources are already assigned to another capacity pool. Some of the resources are directly assigned to a capacity pool, and some are indirectly assigned to a capacity pool at the storage system level. If you continue, the resources will be reassigned to the capacity pool you specified.

BPCUI0258W The following internal resources of a storage system you are attempting to add are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?

Explanation

When a storage system is added to a capacity pool, all of its internal resources (storage pools, NSDs, and file systems) are indirectly assigned to that capacity pool. You are attempting to add one or more storage systems to a capacity pool, but one or more of the storage-system internal resources are already assigned to another capacity pool. If you continue, the resources will be reassigned at the storage system level to the capacity pool you specified.

BPCUI0259W The following storage systems and storage-system internal resources are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?

Explanation

You are attempting to add storage systems to a capacity pool, but one or more of the storage systems and one or more storage-system internal resources (storage pools, NSDs, and file systems) are already assigned to another capacity pool. If you continue, the resources will be reassigned to the capacity pool you specified. At least one of the storage systems you are attempting to add is already assigned to another capacity pool. If you continue, the storage system will be reassigned to the capacity pool you specified. When a storage system is added to a capacity pool, all of its internal resources (storage pools, NSDs, and file systems) are indirectly assigned to that capacity pool. At least one of the storage systems you are attempting to add has one or more internal resources that are already assigned to another capacity pool. If you continue, the resources will be reassigned at the storage system level to the capacity pool you specified.

BPCUI0260E The specified private key file format for the file module is not supported. Please convert it to Open SSH (.pem) key file.

Explanation

Action

BPCUI0261E The service class was not found in the database.

Explanation

This error might occur if the service class was already deleted.

Action

BPCUI0262E The capacity pool was not found in the database.

Explanation

This error might occur if the capacity pool was already deleted.

Action

BPCUI0263E The scheduling of the agent upgrade jobs failed.

Explanation

None of the selected agents were scheduled for upgrade due to errors.

Action

Review the individual detailed error messages.

BPCUI0264E Errors occurred when attempting to schedule the upgrade jobs of some of the agents.

Explanation

Upgrade scheduling of some of the agents was successful, however errors did occur while scheduling some of the upgrades.

Action

Review the individual detailed error messages.

BPCUI0265W Warnings occurred when scheduling the upgrade of *warningCount* of the agents.

Explanation

The agents were successfully scheduled for upgrade, but warnings did occur.

Action

Review the individual detailed warning messages.

BPCUI0266I *attemptedCount* of the selected agents were scheduled for upgrade.

Explanation

All agents that could be upgraded were successfully scheduled for upgrade.

Action

No action.

BPCUI0267I The upgrade agent job was successfully scheduled for *hostName*.

Explanation

Action

No admin response required.

BPCUI0268W Deleting a capacity pool does not affect any volumes or shares that were provisioned from the capacity pool. However, the volumes or shares are no longer associated with the capacity pool. Associations with the following volumes or shares will be removed:

Explanation

This message is for informational purposes only

Action

No action is required.

BPCUI0269W The following volumes are associated with the service class *scName*. When the volumes were created, they satisfied the requirements of the service class. If you modify the service class, the volumes are still associated with the service class, but might not satisfy the new requirements of the service class. Depending on your changes to the service class, users might incorrectly assume that the volumes have properties that they do not possess.

Explanation

BPCUI0270W The following shares are associated with the service class *scName*. When the shares were created, they satisfied the requirements of the service class. If you modify the service class, the shares are still associated with the service class, but might not satisfy the new requirements of the service class. Depending on your changes to the service class, users might incorrectly assume that the shares have properties that they do not possess.

Explanation

BPCUI0271W The following volumes are associated with the service class *scName*. If you delete the service class, the volumes are no longer associated with any service class.

Explanation

BPCUI0272W The following shares are associated with the service class *scName*. If you delete the service class, the shares are no longer associated with any service class.

Explanation

BPCUI0273E The action does not support the specified type of device.

Explanation

An internal operating error occurred. Check the logs for an indication of an error or exception and contact IBM customer support.

BPCUI0274I The connection test to resource *data_Source_Name* was successful.

Explanation

The connection between the server and the target resource was successfully tested.

Action

No action is required.

BPCUI0275I To collect data about zoning or complete zoning actions during provisioning, you must deploy Storage Resource agents to one or more servers that are on the fabric.

Explanation

A Storage Resource agent is required to collect data and complete zoning actions for a fabric.

Action

To deploy Storage Resource agents to a server, go to the Servers > Servers page, click Add Server, and select Deploy an agent for full server monitoring.

BPCUI0276I Agent *agentName* was disabled.

Explanation

The named agent was disabled.

Action

No action.

BPCUI0277I Agent *agentName* was enabled.

Explanation

The named agent was enabled.

Action

No action.

BPCUI0278I The credentials for *agentName* were updated.

Explanation

The credentials of the named device were successfully updated.

Action

No action.

BPCUI0279I There is no job defined for the device *Name*. Please create a job first before running it again.

Explanation

Corresponding job definition for the device cannot be found. Please create a new job definition and try running the job again.

Action

Create a job definition first and attempt the action again.

BPCUI0280I No switches are managed by the *data_Source_Address* data source.

Explanation

The specified data source is not managing any switches. No switch can be detected during the discovery process.

Action

Ensure that information about the data source was entered correctly and that the data source is managing switches. Add the data source again.

BPCUI0282I The resources that are managed by *data_Source_Address* are already known. One or more resources were added.

Explanation

Any resources that are already known are not added again for monitoring. However, one or more new resources were added. The new resources might collect additional information about the managed resources or might collect redundant information about the managed resources.

Action

No further action is required.

BPCUI0284I No fabrics are managed by the *data_Source_Address* data source.

Explanation

The specified data source is not managing any fabrics. No fabric can be detected during the discovery process.

Action

Ensure that information about the data source was entered correctly and that the data source is managing fabrics. Add the data source again.

BPCUI0286I The fabrics that are managed by *data_Source_Address* are already being monitored.

Explanation

You cannot add fabrics that were previously added and configured for monitoring.

Action

To modify a fabric that is being monitored, go to the list page for the fabric. For example, in the navigation pane, select Network Resources > Fabrics. Then, right-click the fabric and select View Properties. In the properties notebook, modify the values for the fabric.

BPCUI0289W The following network shared disks (NSDs) are already assigned to a capacity pool. Are you sure you want to move these NSDs to a different capacity pool?

Explanation

When a file system is added to a capacity pool, all of the NSDs on which the file system resides are indirectly assigned to that capacity pool. You are attempting to add one or more file systems to a capacity pool, but one or more of the related NSDs are already assigned to another capacity pool. If you continue, the NSDs will be reassigned at the file system level to the capacity pool you specified.

BPCUI0290W The following file systems and network shared disks (NSDs) are already assigned to a capacity pool. Are you sure you want to move these resources to a different capacity pool?

Explanation

At least one of the file systems you are attempting to add is already assigned to another capacity pool. If you continue, the file system will be reassigned to the capacity pool you specified.

When a file system is added to a capacity pool, all of the NSDs on which the file system resides are indirectly assigned to that capacity pool. At least one of the file systems you are attempting to add has one or more related NSDs that are already assigned to another capacity pool. If you continue, the resources will be reassigned at the file system level to the capacity pool you specified.

BPCUI0291W The following network shared disks (NSDs) are already assigned to a capacity pool. Are you sure you want to move these NSDs to a different capacity pool?

Explanation

When a storage system is added to a capacity pool, all free NSDs in the storage system are indirectly assigned to that capacity pool. When a file system is added to a capacity pool, all of the NSDs on which the file system resides are indirectly assigned to that capacity pool. You are attempting to add one or more NSDs to a capacity pool, but they are already indirectly assigned to another capacity pool at the storage system level or the file system level. If you continue, the NSDs will be reassigned to the capacity pool you specified. In addition, the storage system or file system will no longer be assigned to its capacity pool. The NSDs that are not being reassigned will instead be directly assigned to the storage system's or file system's original capacity pool.

BPCUI0292E The host name or IP address *ip_address_or_hostname* cannot be reached.

Explanation

This problem might occur if the local area network is down, the resource cannot be reached by a ping (ICMP) command, or a firewall is preventing communication between the server and the resource.

Action

To resolve this problem, try the following actions:

- Verify that the IP address or host name of the resource is valid.
- Verify that the local area network is available.
- Ensure that a firewall is not preventing network access.

Try to add the storage resource again by going to the Storage Systems page in the GUI and selecting Add Storage System.

BPCUI0293I A probe is started for *deviceName*.

Explanation

A data collection job is started to gather probe data for the resource.

BPCUI0294I A performance monitor is started for *deviceName*.

Explanation

A data collection job is started to gather performance data for the resource.

BPCUI0295I The performance monitor is stopped for *deviceName*.

Explanation

The performance monitor is stopped for the resource

BPCUI0297W One resource was added to *capacity_pool_name*. One resource could not be added because it could not be found.

Explanation

One or more resources could not be added to the selected capacity pool because they are no longer in the database.

Action

No action is required.

BPCUI0298W *count* resources were added to *capacity_pool_name*. One resource could not be added because it could not be found.

Explanation

One or more resources could not be added to the selected capacity pool because they are no longer in the database.

Action

No action is required.

BPCUI0299W One resource was added to *capacity_pool_name*. *count_Not_Found* resources could not be added because they could

not be found.

Explanation

One or more resources could not be added to the selected capacity pool because they are no longer in the database.

Action

No action is required.

BPCUI0300W *count* resources were added to *capacity_pool_name*.
count_Not_Found resources could not be added because they could not be found.

Explanation

One or more resources could not be added to the selected capacity pool because they are no longer in the database.

Action

No action is required.

BPCUI0301E Failed to assign the *role name* role.

Explanation

An internal error occurred when trying to assign a role to a list of groups.

BPCUI0302E Failed to retrieve the existing role assignments.

Explanation

An internal error occurred when trying retrieve the existing group to role assignments.

BPCUI0303E Failed to remove all role assignments from the specified groups.

Explanation

An internal error occurred when remove all role assignments from the specified groups.

BPCUI0304W An error occurred when saving the user-defined properties of the *resourceType*.

Explanation

The user-defined properties for one or more resources cannot be saved because of an error on the server. Examples of user-defined properties include the location and display name of the resource. After the resource is added, you can set these properties either on the resource details page or in the properties notebook for the resource.

Action

To set user-defined properties for a resource, go to the resource list page, right-click the resource, and select View Properties. On the properties notebook, click Edit to change the properties and click Save.

For example, to specify the name of a switch, select Network Resources > Switches from the navigation pane. On the Switches page, right-click the switch and select View Properties. On the properties notebook for the switch, click Edit and enter 64 characters or less in the Name field. Click Save.

You can also specify these properties on the details page for a resource. To access the details page, select Network Resources > Switches from the navigation pane, right-click a resource, and select View Details.

BPCUI0305E A capacity pool with the same name already exists.

Explanation

The specified capacity pool name is not unique. A capacity pool with the same name already exists and cannot be duplicated.

Action

Enter a different name for the capacity pool you are want to create.

BPCUI0306W The selected resource was removed, however warnings did occur.

Explanation

The resource was removed however some warnings did occur.

Action

Review the list of detailed messages for further information.

BPCUI0307E The schedule could not be deleted.

Explanation

The schedule could not be deleted.

Action

Check log file for error messages that might help determine the problem.

BPCUI0308I The resource does not have a connection configured. To add a connection to the resource, click Add Storage System.

Explanation

The data source for the resource might have been removed. You must add a connection to the resource again.

Action

Use the Add dialog for the resource to add the connection again.

BPCUI0309I A probe schedule is defined for *deviceName*.

Explanation

A probe schedule is defined for the device.

BPCUI0310I A performance monitor schedule is defined for *deviceName*.

Explanation

A performance monitor schedule is defined for the device.

BPCUI0311I Probe and performance monitor schedules are defined for *deviceName*.

Explanation

Probe and performance monitor schedules are defined for the device.

BPCUI0312I SNMP Discovery of switches is taking longer than expected. Click Close to run the discovery in the background.

Explanation

The process for discovering switches via SNMP is taking longer than expected. This delay might occur because the network connection is slow, the data source is running slowly, or the data source is managing many resources.

Action

To close the window but continue the discovery process in the background, click Close. Try to add the resource again when the discovery is complete.

BPCUI0313I An upgrade is started for server *deviceName*.

Explanation

An upgrade job was started for a server.

BPCUI0314E Failed to retrieve the list of user groups from the WebSphere user repository.

Explanation

An internal error occurred when attempting to retrieve a list of groups from the WebSphere user repository.

BPCUI0315E Failed to retrieve the list of user groups from user repository due to an invalid search string.

Explanation

The search string specified by the user is invalid preventing the query from being submitted to WebSphere.

Action

Specify a valid search string.

BPCUI0316W Failed to update the role cache maintained by the Device server.

Explanation

The role cache maintained by the Device server was not updated.

Action

Restart the Device server.

BPCUI0317E Access can not be removed, because at least one Administrator user must remain in the system.

Explanation

Access can not be removed, because at least one Administrator user must remain in the system.

Action

First add another group having Administrator role and then retry the action.

BPCUI0318E The group mapping can not be modified, because at least one Administrator user must remain in the system.

Explanation

The group mapping can not be modified, because at least one Administrator user must remain in the system.

Action

First add another group having Administrator role and then retry the action.

BPCUI0319I A task is started for resource *resourceName*.

Explanation

A task is started for the specified resource.

BPCUI0320I Probe and performance monitor schedules are defined for *deviceName*. A performance monitor is scheduled to collect performance data after the probe is done.

Explanation

Probe and performance monitor schedules are defined for the device. Performance monitor is disabled and will start after the probe is done.

BPCUI0321I A task is paused for resource *resourceName*.

Explanation

A task is paused for the specified resource.

BPCUI0322E A task could not be paused for resource *resourceName*.

Explanation

A task could not be paused for the specified resource.

BPCUI0323I A task is resumed for resource *resourceName*.

Explanation

A task is resumed for the specified resource.

BPCUI0324E A task could not be resumed for resource *resourceName*.

Explanation

A task could not be resumed for the specified resource.

BPCUI0325E Failed to retrieve the list of users from the WebSphere user repository.

Explanation

An internal error occurred when attempting to retrieve a list of users from the WebSphere user repository.

BPCUI0326E Failed to retrieve the list of users from user repository due to an invalid search string.

Explanation

The search string specified by the user is invalid preventing the query from being submitted to WebSphere.

Action

Specify a valid search string.

BPCUI0327E Failed to get the roles associated with the current user.

Explanation

An internal error occurred when attempting to retrieve the roles associated with the current user.

BPCUI0328I A task is saved.

Explanation

A task is saved.

BPCUI0329I A task was successfully deleted.

Explanation

The task was successfully deleted.

Action

No action is required.

BPCUI0330E The user *user* is not authorized to access the product.

Explanation

The user is not a member of a group that is assigned to a role that is authorized to access the product.

Action

Assign a role to the user's group that provides access to the product.

BPCUI0331I A task is cancelled for resource *resourceName*.

Explanation

A task is cancelled for the specified resource.

BPCUI0332E An unexpected error occurred. The task for schedule *schedule name* could not be paused or resumed.

Explanation

The task type associated with the schedule is not valid.

Action

Refresh the page to check whether the task was completed. If the problem persists, use the Service tool to collect trace data and send it to IBM Software Support.

BPCUI0333E An unexpected error occurred. The task for schedule *schedule name* could not be be paused.

Explanation

The task was not running or the task was completed before the task was paused.

Action

Refresh the page to check whether the task was completed. If the problem persists, use the Service tool to collect trace data and send it to IBM Software Support.

BPCUI0334E An unexpected error occurred. The task for schedule *schedule name* could not be resumed.

Explanation

The task was not paused or the task was completed.

Action

Refresh the page to check whether the task was completed. If the problem persists, use the Service tool to collect trace data and send it to IBM Software Support.

BPCUI0335E The volumes cannot be converted or moved because the target pools do not have sufficient available space.

Explanation

The target pools do not have sufficient available space to complete the operation.

Action

Select target pools with sufficient available space to complete the operation.

BPCUI0336I The ability to provision with block storage devices is only available with the advanced license.

Explanation

The product is installed with the basic license and the functional is unavailable to enable this feature the advance license must be installed.

BPCUI0338E Insufficient user privileges to service the REST request.

Explanation

A request to the REST interface was made by a user who did not have sufficient role privileges.

BPCUI0339E An unexpected error occurred while authorizing the REST request.

Explanation

An internal operating error occurred. Check the logs for an indication of an error or exception and contact IBM customer support.

BPCUI0340I A task was successfully renamed.

Explanation

The task was successfully renamed.

Action

No action is required.

BPCUI0341E The task could not be renamed.

Explanation

The task could not be renamed because the new name already exists.

Action

Enter a different name for the task. Ensure that you have the proper authorization to rename a task.

BPCUI0342E The task could not be renamed because the specified name already exists.

Explanation

The task name that you entered already exists.

Action

Enter a unique name for the task.

BPCUI0343I Performance monitoring is unavailable for resource *resource name* because the resource was not probed.

Explanation

To effectively monitor the performance of a resource, a probe must collect asset and configuration information about that resource.

Action

Schedule a probe to collect information about the resource. After the probe completes, run the performance monitor again.

BPCUI0344W The following service classes allow provisioning only from the *capacity pool* capacity pool: *service classes*. If you delete this capacity pool, the service classes will allow provisioning from any available storage.

Explanation

The capacity pool is the only capacity pool that is associated with one or more service classes. Currently, storage requests that require any of these service classes can be satisfied only by resources in the capacity pool. By removing the capacity pool, you remove this restriction from the service classes.

Action

Identify the reason why each service class specifies the capacity pool restriction. Determine whether removing the restriction is acceptable.

BPCUI0346I The Storage Resource agent that is deployed on the server cannot be uninstalled.

Explanation

The Storage Resource agent cannot be uninstalled.

Action

Ensure that all required servers are running, and that the local area network is available. Verify that you have a network connection to the server that the agent is located on. Wait a few minutes and try again. If none of these actions help resolve the problem, contact IBM Software Support.

BPCUI0347I All servers were removed except for the product server. Entries for the product server resources might still be displayed in the GUI until all the associated removals are complete.

Explanation

You have attempted to remove a number of servers including the Storage Resource agent that is installed on the product server. The agent on the product server cannot be removed.

Action

No further action is required.

BPCUI0348W You cannot provision volumes because at least one of the selected hosts was not found in the database. Ensure that all hosts that are selected for provisioning are being monitored.

Explanation

To request volumes for a server or hypervisor, the server or hypervisor must be added to the database.

Action

Make sure that the server or hypervisor is known. If it is not, add the server or hypervisor using the Add dialog for that type of resource.

BPCUI0349W You cannot provision volumes because not all of the selected hosts appear to have Fibre Channel connectivity.

Explanation

Typically, you cannot request volumes for servers or hypervisors that do not have fabric connectivity. However, if all of the selected hosts have Fibre Channel Port WWPNs and none appear to have fabric connectivity, it is possible that the fabrics were not probed. In this case, you can use the Provision Storage function to request volumes even though the hosts appear to have no fabric connectivity. However, if only some, but not all, of the selected servers or hypervisors appear to have Fibre Channel connectivity, provisioning volumes is not supported.

Action

To request volumes for servers or hypervisors, ensure that either all or none of the selected servers or hypervisors appear to have Fibre Channel connectivity to a managed fabric.

BPCUI0350W You cannot provision volumes because the hypervisors that host the virtual machines use different operating systems.

Explanation

When you provision volumes for one or more virtual machines, the volumes are assigned to the hypervisors that host the virtual machines. To provision a volume to multiple virtual machines, the hypervisors that host the virtual machines must all be running the same operating system type.

Action

To provision volumes to multiple virtual machines, make sure the hypervisors that host the virtual machines all run the same operating system type. The OS Type column on the Hypervisors page displays the operating system type.

BPCUI0351W You cannot provision volumes because there is no Fibre Channel host port information for at least one hypervisor.

Explanation

When you provision volumes for one or more virtual machines that do not have NPIV ports, the volumes are assigned to the hypervisors that host the virtual machines. If you are provisioning to a single virtual machine, the hypervisor that is managing the virtual machine does not have a Fibre Channel host port. If you are provisioning to multiple virtual machines, at least one of the hypervisors that is managing the virtual machines does not have a Fibre Channel host port.

Action

To provision volumes to virtual machines that do not have NPIV ports, ensure that all the hypervisors that are managing the selected virtual machines have Fibre Channel host ports.

BPCUI0352W You cannot provision volumes because not all of the hypervisors that host the virtual machines appear to have Fibre Channel connectivity.

Explanation

When you provision volumes for one or more virtual machines that do not have NPIV ports, the volumes are assigned to the hypervisors that host the virtual machines. Typically, you cannot provision volumes for hypervisors that do not have fabric connectivity. However, if all of the hypervisors that host the virtual machines have Fibre Channel Port WWPNs and none appear to have fabric connectivity, it is possible that the fabrics were not probed. In this case, you can use the Provision Storage function to provision volumes even though the hypervisors appear to have no fabric connectivity. However, if only some, but not all, of the hypervisors that host the virtual machines appear to have Fibre Channel connectivity, provisioning volumes is not supported.

Action

To provision volumes to virtual machines that do not have NPIV ports, ensure that either all or none the hypervisors that are managing the selected virtual machines appear to have Fibre Channel connectivity to a managed fabric.

BPCUI0355W You cannot provision volumes because no block-storage service class exists.

Explanation

When you request volumes by using the Provision Storage wizard, you specify your storage requirements by selecting a block-storage service class. There are currently no block-storage service classes, so you cannot provision volumes.

Action

BPCUI0356W You cannot provision shares because no file-storage service class exists.

Explanation

When you request network-attached storage (NAS) shares by using the Provision Storage wizard, you specify your storage requirements by selecting a file-storage service class. There are currently no file-storage service classes, so you cannot provision shares.

Action

BPCUI0357W You cannot provision volumes because you do not have permission to provision by using any block-storage service class.

Explanation

When you request volumes by using the Provision Storage wizard, you specify your storage requirements by selecting a block-storage service class. If you do not have Administrator privileges, you must have permission to provision by using at least one block-storage service class. You do not have permission to provision by using any block-storage service class, so you cannot provision volumes.

Action

BPCUI0358W You cannot provision shares because you do not have permission to provision by using any file-storage service class.

Explanation

When you request network-attached storage (NAS) shares by using the Provision Storage wizard, you specify your storage requirements by selecting a file-storage service class. If you do not have Administrator privileges, you must have permission to provision by using at least one file-storage service class. You do not have permission to provision by using any file-storage service class, so you cannot provision shares.

Action

BPCUI0359E The credentials for the servers were not updated.

Explanation

None of the credentials for the selected servers were updated.

Action

Review the listed of detailed messages for further information.

BPCUI0360W The credentials for *successfulUpdates* of *attemptedUpdates* of the selected servers were updated.

Explanation

Updating the credentials of the servers was partially successful, but some errors occurred.

Action

Review the listed of detailed messages for further information.

BPCUI0361W The credentials for the selected server was updated, however warnings did occur.

Explanation

The credentials for the selected server was updated, however warnings did occur.

Action

Review the list of detailed messages for further information.

BPCUI0362W The credentials for *successfulUpdates* selected servers were updated, however warnings did occur.

Explanation

The credentials were updated however some warnings did occur.

Action

Review the listed of detailed messages for further information.

BPCUI0363E Cannot connect to the SNMP data source *IP_Address*.

Explanation

A test connection cannot be established to the data source at the specified IP address. This error might occur if the data source is not available or the SNMP community was not entered correctly. The SNMP community name is shared by one or more SNMP hosts and is used to authenticate messages that are received by those hosts..

BPCUI0364I The performance monitor schedule was updated for *deviceName*.

Explanation

A performance monitor schedule was updated for the device.

Action

No further action is required.

BPCUI0366W The server *serverName* was not updated because it does not support the action.

Explanation

The server was not updated because it is an agentless server or it is a server with an agent whose present configuration cannot support the action. Examples of such server actions would be attempting to update the server's credentials when the server is a RXA agent or stopping a server that is already stopped.

BPCUI0367W You cannot provision volumes to virtual machines with NPIV ports and virtual machines without NPIV ports at the same time. To provision volumes to virtual machines, ensure that you select either only virtual machines with NPIV ports or only virtual machines without NPIV ports.

Explanation

By using the Provision Storage wizard, you can provision volumes to multiple virtual machines. You cannot, however, provision to a combination of virtual machines with NPIV ports and virtual machines without NPIV ports. This restriction exists because volumes are assigned directly to a virtual machine only if the virtual machine has an NPIV port. Otherwise, the volume is assigned to the hypervisor that hosts the virtual machine. You cannot provision volumes to servers (including virtual machines) and hypervisors at the same time.

Action

When you open the Provision Storage wizard from the Servers page, make sure that you select only virtual machines with NPIV ports or only virtual machines without NPIV ports.

BPCUI0368W You cannot provision volumes because none of the selected hosts appear to have Fibre Channel connectivity and the automatic zoning option is enabled. Disable the automatic zoning option in your zoning policy.

Explanation

Typically, you cannot request volumes for servers or hypervisors that do not have fabric connectivity. However, because all of the selected hosts have Fibre Channel Port WWPNs and none appear to have fabric connectivity, it is possible that the fabrics were not probed. For this reason, you can use the Provision Storage dialog to request volumes even though the hosts appear to have no fabric connectivity. However, automatic zoning must first be disabled.

Action

Either probe the fabrics or modify the zoning policy to disable automatic zoning. To modify the zoning policy, go to Advanced Analytics - Provisioning - Set Zoning Policy.

BPCUI0369W You cannot provision volumes because none of the hypervisors that manage the selected virtual machines appear to have Fibre Channel connectivity and the automatic zoning option is enabled. Disable the automatic zoning option in your zoning policy.

Explanation

Typically, you cannot request volumes for servers or hypervisors that do not have fabric connectivity. However, because all of the hypervisors that are managing the selected virtual machines have Fibre Channel port WWPNs and none appear to the product to have fabric connectivity, it is possible that the fabrics were not probed. For this reason, you can use the Provision Storage dialog to request volumes even though the hosts appear to have no fabric connectivity. However, you must disable automatic zoning.

Action

Either probe the fabrics or modify the zoning policy to disable automatic zoning. To modify the zoning policy, click Advanced Analytics - Provisioning - Set Zoning Policy.

BPCUI0370E The display name *displayName* is already assigned to resource *resource Name*.

Explanation

The display name of a resource must be unique.

Action

Enter a unique display name for the resource.

BPCUI0372I The selected hosts do not appear to have Fibre Channel connectivity. In the resulting provisioning task, ensure that the recommended storage system is connected to the hosts before you run the task. Also, be aware that all fabric-related options will be ignored.

Explanation

Typically, you cannot request volumes for servers or hypervisors that do not have fabric connectivity. However, because all of the selected hosts have Fibre Channel Port WWPNs and none appear to have fabric connectivity, it is possible that the fabrics were not probed. For this reason, you can use the Provision Storage dialog to request volumes even though the hosts do not appear to have fabric connectivity. However, all of the fabric-related options, such as number of paths, redundant fabrics, and automatic zoning, will be ignored.

Action

If a provisioning task cannot be created, ensure that the hosts are connected to the recommended storage system before you run the task. You must manually configure the fabric. Alternatively, make sure the hosts have fabric connectivity and probe the fabrics.

BPCUI373I Volumes are assigned to the hypervisors that host virtual machines. Volumes are not assigned directly to virtual machines that do not have NPIV ports. None of the hypervisors that manage the virtual machines appear to have Fibre Channel connectivity. In the resulting provisioning task, ensure that the recommended storage system is connected to the hypervisors before you run the task. Also, be aware that all fabric-related options will be ignored.

Explanation

Typically, you cannot request volumes for servers or hypervisors that do not have fabric connectivity. However, because all of the hypervisors that are managing the selected virtual machines have Fibre Channel port WWPNs and none appear to the product to have fabric connectivity, it is possible that the fabrics were not probed. For this reason, you can use the Provision Storage wizard to request volumes even though the hypervisors appear to have no fabric connectivity. However, all of the fabric-related options (number of paths, redundant fabrics, and automatic zoning) will be ignored.

Action

If a provisioning task cannot be created, ensure that the hypervisors are connected to the recommended storage systems before you run the task. You must manually configure the fabric. Alternatively, make sure that the hypervisors have fabric connectivity and probe the fabrics. To assign volumes to virtual machines, wait until the provisioning operation completes. Then, use the hypervisor manager to assign the hypervisor disks to the virtual machines.

BPCUI0374E Schedule is not enabled for the resource *resource*.

Explanation

Schedule is not enabled. Cannot be started.

Action

In order to be started the job should be enabled first.

BPCUI0375E Performance data is not available.

Explanation

Performance data can't be collected, because the Device server is not available.

Action

To check the status of the Device server, complete these steps:

- From the Home menu, click System Management.
- In the navigation pane, click Components.

If the Device server is stopped, restart it.

BPCUI0376E Invalid number of days to keep configuration history. The number should be between *minimum value* and *maximum value*.

Explanation

The input values were not in a valid range.

Action

Retry again with input values within the valid range.

BPCUI0377E Invalid number of days to keep data for removed resources. The number should be between *minimum value* and *maximum value*.

Explanation

The input values were not in a valid range.

Action

Retry again with input values within the valid range.

BPCUI0378E Invalid number of days to keep sample performance data. The number should be between *minimum value* and *maximum value*.

Explanation

The input values were not in a valid range.

Action

Retry again with input values within the valid range.

BPCUI0379E Invalid number of days to keep hourly performance data. The number should be between *minimum value* and *maximum value*.

Explanation

The input values were not in a valid range.

Action

Retry again with input values within the valid range.

BPCUI0380E Invalid number of days to keep daily performance data. The number should be between *minimum value* and *maximum value*.

Explanation

The input values were not in a valid range.

Action

Retry again with input values within the valid range.

BPCUI0381E Failed to update the performance data retention settings.

Explanation

The performance data retention settings on the History Retention page could not be saved.

Action

To resolve the issue, try the following actions:

- Verify that the local area network is available.
- Check the status of the Device server on the Home > System Management page.
- Restart the Device server on the Home > System Management page.
- Verify that that the database repository is up and running.
- Try to save the performance data retention settings again.

BPCUI0382E Performance monitoring is unavailable for resource *resource name*.

Explanation

Performance monitoring is not supported for the device.

Action

Performance monitoring is not supported for the device.

BPCUI0383E Failed to update the history retention settings.

Explanation

The history retention settings on the History Retention page could not be saved.

Action

To resolve the issue, try the following actions:

- Verify that the local area network is available.
- Check the status of the Device server on the Home > System Management page.
- Restart the Device server on the Home > System Management page.
- Verify that that the database repository is up and running.
- Try to save the history retention settings again.

BPCUI0384E Failed to retrieve the history retention settings.

Explanation

The most recent history retention settings could not be displayed on the History Retention page.

Action

To resolve the issue, try the following actions:

- Verify that the local area network is available.
- Check the status of the Device server on the Home > System Management page.
- Restart the Device server on the Home > System Management page.
- Verify that the database repository is up and running.
- Reload the History Retention page.

BPCUI0385E Invalid number of runs to keep log files for each schedule. The number should be between *minimum value* and *maximum value*.

Explanation

The input values were not in a valid range.

Action

Retry again with input values within the valid range.

BPCUI0386E A job cannot be run for resource *resourceName* because there is a job already running for the resource. Wait for the job to finish and try again.

Explanation

A new job cannot be run for a schedule definition if a previous job from that schedule definition did not finish running. Only one running job per schedule at a time is permitted.

Action

BPCUI0387I The selected resources support different performance monitor intervals. If you select multiple resources, intervals that are common to all resources are displayed in the interval list.

Explanation

The selected resources do not support the same set of performance monitor intervals. For example, you select two storage systems. Storage_System_a supports performance monitor intervals of 5 minutes, 10 minutes, and 15 minutes. Storage_System_b supports intervals of 10 minutes, 15 minutes, and 20 minutes. The interval list displays intervals of 10 minutes and 15 minutes.

Action

To configure a performance monitor with an interval that is not in the interval list, you can schedule the performance monitor for a single resource or a group of resources of the same type.

BPCUI0388E The probe schedule cannot be created for resource {0} because not all the information was provided. If you are configuring a probe for a resource for the first time, you must enter values for the probe status, time, and frequency fields.

Explanation

You attempted to create a probe schedule but not all the information was provided. If you are configuring a probe for a resource for the first time, you must enter values for the probe status, time, and frequency fields.

Action

When you create a probe schedule, enter all the information that is required. If you are modifying probe schedules for multiple resources at the same time, do not include resources that are not configured for probe data collection.

BPCUI0389E The performance monitor schedule cannot be created because not all the information was provided. If you are configuring a performance monitor for a resource for the first time, you must enter values for the performance monitor status and interval fields.

Explanation

You attempted to create a performance monitor but not all the information was provided. If you are configuring a performance monitor for a resource for the first time, you must enter values for the performance monitor status and interval fields.

Action

When you create a performance monitor schedule, enter all the information that is required. If you are modifying performance monitors schedules for multiple resources at the same time, do not include resources that are not configured for performance monitor data collection.

BPCUI0390I The service logs were successfully created.

Explanation

The service logs zip file were successfully created on the server.

Action

The user can either download the logs by clicking the link on the message or they can download later from the System Management Overview page.

BPCUI0391I The connection test to data source *data source* was successful. A probe is running. The health status is unknown until the probe is finished.

Explanation

The previous state of the monitored device was either unreachable or unknown. The health status of the device is unknown until the currently running probe updates the status.

Action

No action is required.

BPCUI0392I The connection test to the data source *data source* was successful.

Explanation

The connection test was successful. However, to view the most recent information about the resource, you must run a probe.

Action

BPCUI0393E The user *user_name* does not have sufficient privileges to deploy the vSphere Web Client extensionr.

Explanation

The user must have Administrator privileges to deploy the vSphere Web Client extension.

Action

Specify a vSphere user name with Administrator privileges.

BPCUI0394E The user *user_name* does not have permission to log in to the vCenter Server system.

Explanation

The user does not have permission to log in to the vCenter Server system.

Action

Specify a user with permission to log in to the vCenter Server system.

BPCUI0395E This version of the vCenter Server *server_name* does not support the deployment of the vSphere Web Client extension for the product.

Explanation

This version of vCenter Server does not support the deployment of the vSphere Web Client extension.

Action

Deploy the required vSphere Web Client extension on vCenter Server version 5.1 or later.

BPCUI0396E The user *user_ID* does not have the required role. The role associated with this user must be Administrator, Monitor, or External Application.

Explanation

To access the vSphere Web Client extension and VASA functionality, the user ID must have one of the following roles: Administrator, Monitor, or External Application.

Action

Specify a user ID with one of the required roles.

BPCUI0397E The vCenter Server user name or password is invalid.

Explanation

The user name or password is not valid.

Action

Ensure that you enter the correct user name and password.

BPCUI0398E The user name or password is invalid.

Explanation

The user name or password is not valid.

Action

Ensure that you enter the correct user name and password.

BPCUI0399I The server was started.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCUI0400E Failed to retrieve the system management information from the Data server.

Explanation

An error occurred when attempting to retrieve the system management information from the Data server.

Action

Verify that the Data server is running.

BPCUI0402E Failed to retrieve the server status of the Data server.

Explanation

An error occurred when attempting to retrieve the server status of the Data server.

Action

Verify that the Data server is running.

BPCUI0403E The SMI-S provider service is not available.

Explanation

The service that is accessed by the SMI-S provider is not available.

Action

Make sure that the service accessed by the SMI-S provider is started and operational.

BPCUI0404E An error occurred while updating the trace log configuration file. The original file *file* was deleted and could

not be restored. A backup of this file may be available at *backup file*.

Explanation

An error occurred during the process of updating the Performance Management trace log configuration file. The original file was deleted and needs to be restored. This will not impact the current operation of the Performance Manager but the file is needed prior to restarting the Device server.

Action

Restore the backup of the Performance Management configuration file, if possible. Otherwise contact IBM support.

Related reference

- [Getting support](#)

BPCUI0405E Failed to set the trace settings from the Data server.

Explanation

An error occurred when attempting to save the trace settings from the Data server.

Action

Verify that the Data server is running.

BPCUI0406E Cannot start the server. The start script reported the following error: *error*

Explanation

An error occurred when executing the start script for the server.

Action

Retry the operation or manually start the server.

BPCUI0407E Cannot start the server. Unable to locate the start script *path to script*.

Explanation

The script required to start the server could not be found on the local filesystem.

Action

Verify that the script exists at the specified location.

BPCUI0408E Cannot start the server. Unable to execute the start script *path to script*.

Explanation

The script required to start the server could not be executed.

Action

Verify that the script is configured with the required file permissions so that it can be run by the product.

BPCUI0409W The server is taking a long time to start. If the server status continues to show an error status after a reasonable interval, try to start the server again.

Explanation

The server is in the process of starting but it is taking a long time.

Action

To resolve this issue, try the following actions:

1. Wait for the Alert server to start. If the Alert server does not start after a reasonable amount of time, try restarting the server.
2. Make sure the database is running properly.
3. Make sure the network is functioning properly.
4. Make sure the server or file system has not run out of disk space.
5. In some rare cases, the following directories or files might be missing or corrupt:
 - (installation_directory)/alert/conf directory
 - (installation_directory)/wlp/usr/servers/alertServerjvm.optionsIf this is the case, restore these directories and files from the installation image.

BPCUI0410E Cannot stop the server. The stop script reported the following error: *error*

Explanation

An error occurred when executing the stop script for the server.

Action

Retry the operation or manually stop the server.

BPCUI0411W The server is taking a long time to stop. If the server status continues to show that it is still running try to stop the server again after a reasonable interval.

Explanation

The server is in the process of stopping but it is taking a long time to stop.

Action

BPCUI0412E Cannot stop the server. Unable to locate the stop script *path to script*.

Explanation

The script required to stop the server could not be found on the local filesystem.

Action

Verify that the script exists at the specified location.

BPCUI0413E Cannot stop the server. Unable to execute the stop script *path to script*.

Explanation

The script required to stop the server could not be executed.

Action

Verify that the script is configured with the required file permissions so that it can be run by the product.

BPCUI0414W It is taking a long time for the services to start. If the server status continues to show an error status after a reasonable interval, try to start the services again. If the problem persists then restart the server.

Explanation

The request to start the services has been issues but it is taking a long time for them to start.

Action

BPCUI0415E Failed to start the service *service name*.

Explanation

An error occurred when attempting to start the specified service of the Data server.

Action

Retry to start the service. If the error continues to occur try restarting the Data server.

BPCUI0416I The server was stopped.

Explanation

This message is for informational purposes only.

Action

No action is required.

BPCUI0417I The services of the server were started.

Explanation

The user initiated the starting of internal server services that were not running

Action

The services have been started, no further action is required.

BPCUI0418E The action cannot be completed because the data source that is managing this resource cannot be reached.

Explanation

A test connection cannot be established to the data source that is managing this resource.

Action

Ensure that the data source is available. Ensure that the network is available and that you have a network connection to the data source. Verify that a firewall is not preventing network access to the data source. Try the action again. If the problem persists, try adding the data source again.

BPCUI0419E A Storage Resource agent is already deployed for this server and has a status of Pending deployment or Failed deployment. Use the Servers page to resolve the deployment errors or modify the deployment schedule.

Explanation

The Storage Resource agent is already added. Use the Servers page to check the status of the agent deployment.

Action

On the Servers page, if the server has a status of Pending deployment, you can cancel the deployment or modify the deployment schedule. If the server has a status of Failed deployment, you can cancel the deployment or resolve the errors that caused the deployment to fail and then deploy the agent again.

BPCUI0420E A file access error occurred when the system attempted to back up or modify the tracing configuration file *configuration file*.

Explanation

The action to change tracing levels first backs up, and then modifies, the tracing configuration file. The error occurred during that process.

Action

Review the trace file for the exception details. Look at the directories and file information to determine why the exception occurred.

BPCUI0421E There is a log collection operation already running. A new one cannot be submitted until the current one completes.

Explanation

Only one log collection operation can run on the server at a time. There is already one running.

Action

BPCUI0422E Cannot start the log collecting job. Unable to locate the required script *path to script*.

Explanation

The script that is required to collect the logs was not found on the local file system.

Action

Verify that the script exists at the specified location.

BPCUI0423E Cannot start the log collecting job. Unable to run the log collection script *path_to_script*.

Explanation

The script that collects the product logs could not be run.

Action

Verify that the script is configured with the correct permissions to allow it to be run from the web server.

BPCUI0424E Storage cannot be provisioned from capacity pool *capacity pool* using service class *service class* for the following reason:

Explanation

An error occurred during the execution of the provisioning plan for the specified service class and capacity pool.

Action

Verify that the capacity pool for the service class still meets the requirements for the requested storage to be provisioned.

BPCUI0425W The task *task name* cannot be scheduled because it is already running.

Explanation

This task cannot be scheduled to run later because it is already running. It cannot be run again.

Action

No action is required.

BPCUI0426E Storage cannot be provisioned by using service class *service class* for the following reason:

Explanation

A provisioning task cannot be created. The requested capacity cannot be found for the selected service class.

Action

Verify that there is free space in storage resources that meet the requirements of the selected service class.

BPCUI0427W The selected group action is complete for all tasks, but warnings were reported.

Explanation

The action to execute, schedule, or delete tasks is complete. However, the action resulted in at least one warning for one or more of the tasks.

Action

The warnings are listed following this message. Refer to the warning messages.

BPCUI0428I The selected group action is complete for all tasks. Some informational messages were returned.

Explanation

The action to execute, schedule, or delete tasks was successful. No warnings or errors were reported. At least one informational message was returned for one or more of the tasks.

Action

No action is required.

BPCUI0429E The validation process cannot contact the server. The server might be down or unreachable due to network problems.

Explanation

The server cannot be added because the validation process cannot contact the server.

Action

Use the ping command to verify that you have a network connection to the server. Try to add the server again.

BPCUI0430I Some tasks were not deleted because they were already run.

Explanation

When you complete the steps of the Provision Storage wizard, provisioning tasks are created. If you return to the Provision Storage wizard without executing, scheduling, or deleting the provisioning tasks, the tasks are deleted. Not all the tasks were deleted, however, because some were already executed.

Action

No action is required.

BPCUI0431E Failed to retrieve the list of managed devices.

Explanation

An unexpected error occurred trying to retrieve the list of managed devices from the database.

Action

Check the connection to the database and retry the operation.

BPCUI0432E Failed to retrieve the performance monitoring granularity from the Device server. Check the connection to the Device server and retry the operation.

Explanation

An unexpected error occurred trying to retrieve the performance monitoring granularity from the Device server.

Action

Check the connection to the Device server and retry the operation.

BPCUI0433E OS type *osType* specified on line *line* of file *file* is not valid.

Explanation

The specified OS type in the file is not valid.

Action

Enter a valid OS type. Valid OS types are "Windows", "Linux", "AIX", "Solaris", or "HP-UX".

BPCUI0434E Data source *data_Source_Key* could not be found.

Explanation

The data source could not be found in the database.

Action

Ensure that the database is running properly. Ensure that the Data server and Device server are up and running. Verify that the local area network is available and a firewall is not preventing network access to product services and agents. Try the action again.

If the problem persists, check the log files for error messages that might help determine the problem. For information about the location of log files and how to start the Data server, Device server, and Db2 database repository, see the IBM Knowledge Center.

If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCUI0435E Required host name or IP address and OS type were not specified on line *line* of file *file*.

Explanation

Deploying a server from a file requires a minimum host name or IP address and OS type.

Action

Ensure the file contains the required host name or IP address and OS type.

BPCUI0436E The alert notification settings cannot be displayed.

Explanation

Information about the alert notification settings for the resource cannot be retrieved from the product database.

Action

Ensure that the product database, Data server, and Device server are up and running. Verify that the local area network is available. Try to view the alert notification settings again.

BPCUI0437E The alert notification settings cannot be saved.

Explanation

Information about the alert notification settings for the resource cannot be saved to the product database.

Action

Check the log files for the Data server and the GUI to determine where the problem might have occurred. The log files are located in the following default locations:

- Data server: (Windows) TPC_installation_directory\data\log, (Linux or AIX) TPC_installation_directory/data/log
- GUI: (Windows) TPC_INSTALL_DIR\web\log, (Linux or AIX) TPC_INSTALL_DIR/web/log

Ensure that the product database, Data server, and Device server are up and running. Verify that the local area network is available. Try to save the alert notification settings again.

BPCUI0438E File *file* does not exist or is empty.

Explanation

The script file that you want to upload must exist in the specified location and contain the commands that you want to run when an alert is detected.

Action

Ensure that the script file exists in the specified location. If it does not exist, create the file or specify a different location. If the file exists but is empty, edit the file to include the script commands that you want to run when an alert is detected.

BPCUI0439E The file *file* could not be uploaded.

Explanation

An error occurred while uploading the script

Action

Check the log files for the GUI to determine where the problem might have occurred. The log files are located in the following default locations:

- GUI: (Windows) TPC_INSTALL_DIR\web\log, (Linux or AIX) TPC_INSTALL_DIR/web/log

BPCUI0440E The text *location* specified on line *line* of file *file* has invalid character(s): *characters*

Explanation

The specified location or custom tag in the file contains invalid characters.

Action

Enter a valid location or custom tag that must contain only letters, numbers, spaces and common characters.

BPCUI0441E The alert definitions cannot be displayed.

Explanation

Information about the alert definitions for the resource cannot be retrieved from the product database.

Action

Ensure that the product database, Data server, and Device server are up and running. Verify that the local area network is available. Try to view the alert definitions again.

BPCUI0442E The alert definitions cannot be saved.

Explanation

Information about the alert definitions for the resource cannot be saved to the product database.

Action

Check the log files for the Data server and the GUI to determine where the problem might have occurred. The log files are located in the following default locations:

- Data server: (Windows) TPC_installation_directory\data\log, (Linux or AIX) TPC_installation_directory/data/log
- GUI: (Windows) TPC_INSTALL_DIR\web\log, (Linux or AIX) TPC_INSTALL_DIR/web/log

Ensure that the product database, Data server, and Device server are up and running. Verify that the local area network is available. Try to save the alert definitions again.

BPCUI0443E Select at least one managed server that is deployed for which alert notification settings need to be displayed.

Explanation

Alert notification settings can be displayed only for managed servers that are deployed.

Action

Select at least one managed server that is deployed and try to view the alert notification settings again.

BPCUI0444E Select at least one managed server that is deployed for which alert definitions need to be displayed.

Explanation

Alert definitions can be displayed only for managed servers that are deployed.

Action

Select at least one managed server that is deployed and try to view the alert definitions again.

BPCUI0445W The discovery job completed with errors. Some available devices were not discovered.

Explanation

The discovery job was able to discover some devices, but encountered problems discovering other available devices.

Action

Check the logs for an explanation of errors. Correct the errors and run the discovery again, or continue configuring the devices that were discovered.

BPCUI0446E Unable to test the connection to the device because the request was not processed by the data collector.

Explanation

The data collector did not respond to the server in the allotted time. The data collector might not be running or it might not be able to connect to the server.

Action

Verify that the data collector is running and that it can connect to the server.

BPCUI0447E Select at least one managed storage subsystem for which alert notification settings need to be displayed.

Explanation

Alert notification settings can be displayed only for managed storage subsystems.

Action

Select at least one managed storage subsystem and try to view the alert notification settings again.

BPCUI0448E Select at least one managed storage subsystem for which alert definitions need to be displayed.

Explanation

Alert definitions can be displayed only for managed storage subsystems.

Action

Select at least one managed server that is deployed and try to view the alert definitions again.

BPCUI0449E The user does not have the required authority to complete the task or command.

Explanation

The level of authority that is required for the task or command depends on the type of resource that the user is managing. This error can occur in the following situations:

- For IBM SONAS resources, if the user is an rssh restricted account and can issue only a restricted set of commands on the resource.
- For IBM Spectrum Scale, if the user that is used to log on to the cluster node does not have privileges to monitor the GPFS cluster. The user must have root privileges on the cluster node or have privileges to run a set of specified administration commands using the sudo command. For information about monitoring IBM Spectrum Scale without requiring root privileges, go to the IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/search/tpch_t_configuring_sudo_access?scope=SS5R93. View the appropriate topic for the version of IBM Spectrum Control that you are using.

Action

Ensure that the user has the required permissions to complete the task.

BPCUI0451E One or more applications from provided list: *names* do not exist.

Explanation

An application with the specified name was not found.

Action

Verify that the required application name exists. If it does not, specify a different application.

BPCUI0452E *entity name* is not supporting data collection actions.

Explanation

The specified entity does not support data collection actions.

Action

The specified entity does not support data collection actions.

BPCUI0453E One or more departments from provided list: *names* do not exist.

Explanation

A department with the specified name was not found.

Action

Verify that the required department name exists. If it does not, specify a different department.

BPCUI0455I No performance data is available for the selected resources.

Explanation

Performance data must be collected about a resource before you can view its performance. Use performance monitors to collect performance data about a resource.

Action

To collect performance data about a resource immediately, go to the Storage Systems page or Switches page, right-click the storage system or switch, and select Data Collection > Start Performance Monitor. To schedule performance monitors to run at set times, select Data Collection > Schedule.

BPCUI0456E You cannot complete the action because the service is temporarily unavailable.

Explanation

IBM Marketplace is investigating the issue and service will be resumed as soon as possible.

Action

Wait a few minutes and try again. If you still can't complete the action, go to your Products and services page (<https://myibm.ibm.com/products-services/>) on IBM Marketplace. Click the down-arrow for the Storage Insights offering, click Support, and then choose an option.

BPCUI0457W The applications *listOfApplications* cannot be deleted because they contain subcomponents *subcomponent*, which cannot be moved up a level in the applications hierarchy due to name conflicts with existing applications in that higher level.

Explanation

When you delete an application, the subcomponents that belong to that application are not automatically deleted unless you specify the Remove subcomponents option. They are moved to the next higher level in the application hierarchy. If those subcomponents have the same name as an application in the higher level, they cannot be moved up, and the parent application of those subcomponents cannot be deleted.

Action

Either rename the specified subcomponents and then delete the applications, or specify the Remove subcomponents option when you delete these applications.

BPCUI0458W The departments *listOfDepartments* cannot be deleted because they contain subdepartments or applications *subdepartment*, which cannot be moved up a level in the departments hierarchy due to name conflicts with departments in that higher level.

Explanation

When you delete a department, the subdepartments and applications that belong to that department are not automatically deleted unless you specify the Remove subdepartments and applications option. They are moved to the next higher level in the department hierarchy. If those subdepartments and applications have the same name as a department or an application in the higher level, they cannot be moved up, and the parent department of those subdepartments and applications cannot be deleted.

Action

Either rename the specified subdepartments or applications and then delete the department, or specify the Remove subdepartments and applications option when you delete these departments.

BPCUI0459W The selected subcomponents cannot be removed from the application because they cannot be moved up a level in the application hierarchy due to name conflicts with the existing applications or subcomponents at the higher level.

Explanation

When you remove a subcomponent from an application, the subcomponent is moved up to the same level as the parent application. If the subcomponent has the same name as an existing application or subcomponent at that higher level, it cannot be moved up and the subcomponent is not removed from the application.

Action

Rename the specified subcomponents and try to remove them from the application again.

BPCUI0460W The selected applications or subdepartments cannot be removed from the department because they cannot be moved up a level in the department hierarchy due to name conflicts with the existing applications or subdepartments at the higher level.

Explanation

When you remove an application or subdepartment from a department, the application or subcomponent is moved up to the same level as the parent department. If the application or subdepartment has the same name as an existing application or subdepartment at the higher level, it cannot be moved up. The application or subdepartment is not removed from the department.

Action

Rename the specified applications or subdepartments and try to remove them from the department again.

BPCUI0461W There are no task details to display. The analysis-execution task could not be run.

Explanation

The execution task was created when the Device server was not available. The task is not valid.

Action

Ensure that the Device server is running, and then create and run a new task.

BPCUI0462E Failed to add the device because the data collector is not responding.

Explanation

The data collector did not respond to the server in the allotted time. The data collector might not be running or it might not be able to connect to the server.

Action

Verify that the data collector is running and that it can connect to the server.

BPCUI0463E The discovery failed because the data collector is not responding.

Explanation

The data collector did not respond to the server in the allotted time. The data collector might not be running or it might not be able to connect to the server.

Action

Verify that the data collector is running and that it can connect to the server.

BPCUI0464E The connection test failed because the data collector is not responding.

Explanation

The data collector did not respond to the server in the allotted time. The data collector might not be running or it might not be able to connect to the server.

Action

Verify that the data collector is running and that it can connect to the server.

BPCUI0465E The requested action failed because the data collector is not responding.

Explanation

The data collector did not respond to the server in the allotted time. The data collector might not be running or it might not be able to connect to the server.

Action

Verify that the data collector is running and that it can connect to the server.

BPCUI0466I The servers were created.

Explanation

Action

The action completed successfully, no further action necessary

BPCUI0467W *successCount* of *totalCount* servers were created.

Explanation

Action

Not all of the servers were created, review the detailed error messages for more information.

BPCUI0468E The creation of the servers failed.

Explanation

Action

None of the servers were created, review the detailed error messages for more information.

BPCUI0469E Schedule job does not exist for *entity name*.

Explanation

Schedule job should be created first in order to be started.

Action

Schedule job should be created first in order to be started.

BPCUI0470E Invalid file *file* size of *size* GB. Maximum allowed file size is *max size* GB.

Explanation

Maximum allowed files size exceeded.

Action

Verify the size of the input file.

BPCUI0471E Failed to set the trace settings from the Alert server.

Explanation

An error occurred when attempting to save the trace settings from the Alert server.

Action

Verify that the Alert server is running.

BPCUI0472E Failed to retrieve the system management information from the Alert server.

Explanation

An error occurred when attempting to retrieve the system management information from the Alert server.

Action

Verify that the Alert server is running.

BPCUI0474E Failed to retrieve the server status of the Alert server.

Explanation

An error occurred when attempting to retrieve the server status of the Alert server.

Action

Verify that the Alert server is running.

BPCUI0475I The volumes have been excluded from the reclamation analysis.

Explanation

The removed volumes will now be in the excluded list and will not be considered in future runs of the reclamation analysis.

Action

The action completed successfully, no further action necessary.

BPCUI0476I The volumes will be included in future analyses to reclaim storage.

Explanation

On completion of the next reclamation analysis the volumes may appear in the table of recommended reclamation volumes and reclamation summary charts

Action

The action completed successfully, no further action necessary.

BPCUI0477E An unexpected error occurred when modifying the optimization characteristics of the volumes.

Explanation

To find the cause of the issue, further investigation is required.

Action

To resolve the issue, try the following actions:

- Check the status of the product servers on the Home > System Management page.
- Verify that the database repository is available.
- Verify that the related database service is active.
- Check for error messages in the log files for the servers.

If you need more information, go to IBM Knowledge Center and check the Administering section. You can access IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93>. If you cannot resolve the issue, contact IBM Support.

BPCUI0478E The scheduled agent upgrade time is in the past.

Explanation

The upgrade was scheduled to occur at a time in the past.

Action

Schedule another date and time to upgrade the agent.

BPCUI0479E The object storage credentials are incorrect. Enter the correct credentials. Alternatively, clear the object credentials check box and do not specify the authentication credentials for object storage now. You can use the Modify Connection action to add the object storage later.

Explanation

The GPFS cluster is configured for object storage, but the specified user credentials are invalid.

Action

Enter the correct object storage credentials and try the operation again.

Alternatively, clear the object credentials check box and do not specify the authentication credentials for object storage now. You can use the Modify Connection action on the Object Storage Systems page to add the object storage later.

BPCUI0480E An object storage request failed on the GPFS cluster.

Explanation

The IBM Spectrum Control server cannot connect to the object storage on the GPFS cluster. This error might occur because the object service is disabled or stopped.

Action

Verify that the object service is configured correctly and is enabled and started. For more information about configuring the object service for IBM Spectrum Scale, go to the IBM Knowledge Center (<http://www.ibm.com/support/knowledgecenter>).

Try the operation again. If the problem persists, contact IBM Software Support.

BPCUI0481W No resources were removed.

Explanation

None of the selected resources were removed, either the resource was already removed or could not be found.

Action

Wait a few minutes, if the resource is still present contact IBM Support.

BPCUI0482E No resources were updated.

Explanation

None of the selected resources were updated because the resource could not be found.

Action

Contact IBM Support.

BPCUI0483E The connection information cannot be updated because it points to another device.

Explanation

The serial number of the device that is managed by this data source doesn't match the serial number of the existing device.

Action

Enter the host name or the IP address of the existing device.

BPCUI0484I The connection information for *device name* was updated.

Explanation

You successfully updated the connection information of the named device.

Action

No action is required.

BPCUI0485E The connection information cannot be updated.

Explanation

The connection information cannot be updated because the serial number of the resource cannot be retrieved. This problem might occur if the resource is unavailable or the connection to the resource was lost.

Action

Try the following actions:

- Verify that the resource is available.
- Ensure that a connection to the resource is active.
- Verify the available local disk space.
- Check the log files for error messages to determine the cause. For information about the location of log files, see the IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93/>.

BPCUI0486E Cannot query the object service for information about accounts and containers as the specified user does not have admin privileges.

Explanation

To query account and container information from the object service, the user must be assigned the admin role in Keystone, the OpenStack identity service. To monitor all accounts and containers, the user must also be assigned the role that is defined in the `reseller_admin_role` configuration option in the Swift proxy server. The default value for the `reseller_admin_role` option is `ResellerAdmin`.

Action

Check the credentials for the user. In Keystone, the OpenStack identity service, ensure that the user is assigned the required role and has the authority to retrieve object storage account and container information. Try the operation again.

BPCUI0487I The connection information of the selected device was successfully updated. Other devices were detected as being managed by the same data source. Would you like to update the connection information of all of them?

Explanation

Other devices were found in Spectrum Control as being managed by the same data source.

Action

Select "Yes" to update the connection information of all detected devices, or "No" to close the dialog.

BPCUI0488I The connection information of all devices connecting through this data source was updated.

Explanation

You successfully updated the connection information of all devices connecting through this data source.

Action

No action is required.

BPCUI0489W Some of the devices connecting through this data source failed to be updated.

Explanation

One or more devices connecting through this data source was not updated.

Action

See below the reason of failure for each device.

BPCUI0490I The *vCenter* vCenter Server was removed.

Explanation

The removal of the specified vCenter Server succeeded.

Action

No action is required.

BPCUI0491E The *vCenter* vCenter Server was not found in the database.

Explanation

The specified vCenter Server was not found in the database. Maybe it was removed.

Action

Wait a few minutes, then verify that the vCenter Server is not listed in the "vCenters" tab of the "Hypervisors" page.

BPCUI0492E The *selected vCenter Servers* were not found in the database.

Explanation

None of the selected vCenter Servers was found in the database. Maybe they were removed.

Action

Wait a few minutes, then verify that the vCenter Servers are not listed in the "vCenters" tab of the "Hypervisors" page.

BPCUI0493I The *vCenter* vCenter Server and all *number of monitored hypervisors* hypervisors monitored by it were successfully removed.

Explanation

The selected vCenter Server and the hypervisors that it monitors were successfully removed.

Action

No action is required.

BPCUI0494I The *number of vCenters* selected vCenter Servers and all *number of monitored hypervisors* hypervisors monitored by them were successfully removed.

Explanation

The selected vCenter Servers and the hypervisors that they monitor were successfully removed.

Action

No action is required.

BPCUI0495W Only *number of removed vCenters* of *number of selected vCenters* of the selected vCenter Servers and *number of removed monitored hypervisors* of *number of monitored hypervisors* of the hypervisors monitored by them were successfully removed.

Explanation

The selected vCenter Servers and the hypervisors that they monitor were partially removed, but some errors occurred.

Action

Review the list of detailed messages for further information.

BPCUI0496I The following fabrics were detected as being managed by the same data source: *comma separated fabrics list*. This action applies to all fabrics that are managed by the current data source. Would you like to update the connection information of all of them?

Explanation

Other monitored fabrics are being managed by the same data source. When you change the connection information for one fabric, the change will apply to all of the fabrics.

Action

Select Yes to update the connection information of all monitored fabrics, or No to close the window without updating connection information.

BPCUI0497E The following fabrics cannot be monitored through the SMI agent: *comma separated fabrics list*. The data source connection information will not be updated.

Explanation

This message applies to all fabrics that are managed by the SMI agent and occurs when one or more switches in the fabrics cannot be reached through the SMI agent.

Action

To resolve the problem, try the following actions:

- Ensure that all of the switches in the fabrics can be reached by the SMI agent.
- Use a different SMI agent to connect to a fabric. To connect a fabric to a different SMI agent, go to Network > Fabrics, click Add Fabric, and select Other. Specify information about the SMI agent, including host name or IP address, to complete the add action.

BPCUI0498E The fabric cannot cannot be monitored through the SMI agent.

Explanation

At least one switch in this fabric cannot be monitored through the SMI agent.

Action

Ensure that all of the switches in the fabric can be reached by the SMI agent.

BPCUI0499I Other switches were detected as being managed by the same data source. This action applies to all switches that are managed by the current data source. Would you like to update the connection information of all of them?

Explanation

Other monitored switches are being managed by the same data source. When you change the connection information for one switch, the change will apply to all of the switches.

Action

Select Yes to update the connection information of all monitored switches, or No to close the window without updating connection information.

BPCUI0500E One or more switches cannot be monitored through the SMI agent. The data source connection information will not be updated.

Explanation

This message applies to all switches that are managed by the SMI agent and occurs when at least one switch cannot be reached through the SMI agent.

Action

Ensure that all of the switches can be reached by the SMI agent.

BPCUI0501E The information cannot be displayed. Log out of the GUI, log in, and try the action again.

Explanation

An error occurred when the server tried to process a request. This message might occur when a request is not handled correctly by the internal Web server, the Web server is down, or the related data collection job failed.

Action

To resolve the problem, try the following actions:

- Ensure that the Web server (and all required product servers) are running and that you have a network connection to the system on which they are located. To check the status of the servers, go to the Home > System Management page. If the status of any server is not Running, click Component Servers for more details.
- Ensure that the related data collection jobs are running and completing successfully. To check the data collection jobs for a specific resource, go to the resource page (for example, Storage > Block Storage Systems) and check the Probe Status and Performance Monitor status columns.
- Log out of the GUI, log in, and try the action again.
- Restart the browser, log in to the GUI, and try the action again.

If the problem persists, contact IBM support. For information about how to contact IBM support, go to <http://www.ibm.com/support/knowledgecenter/SS5R93> and navigate to Troubleshooting and problem determination > Contacting IBM Software Support.

BPCUI0502E The device is already managed by this data source. The data source connection information will not be updated.

Explanation

The device is already being managed by the data source.

Action

No action is required.

BPCUI0503I The connection information of the selected switches was updated.

Explanation

The selected switches have been updated with new connection information.

Action

No action is required.

BPCUI0504I The detected versions of the resources discovered on the data source *data_Source_Address* are unsupported.

Explanation

This resource is of an unsupported version.

Action

For a list of supported versions of this product, go to the support matrix at <http://www.ibm.com/support/docview.wss?uid=swg21386446>.

BPCUI0505E The resource does not have a connection configured.

Explanation

The data source for the resource might have been removed. You must add a connection to the resource again.

Action

To add the connection to the resource again, go to the resource list page for this resource type and click the Add button.

BPCUI0506E Cannot connect to the Alert server.

Explanation

The Alert server is unavailable. This error might occur if the Alert server is down or the local area network is unavailable.

Action

Verify that the database service and Alert server are up and running. Ensure that you have a network connection to the server on which the Alert server is located. Try the action again.

To check the status of IBM Spectrum Control servers and services, go to Home - System Management. For additional status details, click Component Servers or Database in the left navigation pane.

BPCUI0507E The version of the *tpc_server* IBM Spectrum Control Server is not supported.

Explanation

You can manage only the supported versions of IBM Spectrum Control.

Action

Check the product support site for a list of IBM Spectrum Control Server versions that are supported.

BPCUI0508E Cannot connect to the rollup server *rollup_server* on port *host_port*.

Explanation

The rollup server cannot be contacted on the specified port.

Action

To resolve the issue, try the following actions:

- Ensure that you enter the correct host address and port.
- Verify that the local area network is available and a firewall is not preventing network access to product services.
- Check the status of the rollup server.
- Ensure that the Device server is running properly.
- Ensure that the database is running properly.
- Verify that the database repository is available.
- Verify that the related database service is active.
- Check for error messages in the server log files.

BPCUI0509E Cannot authenticate with the rollup server using the provided credentials.

Explanation

The user name or password that was entered for the rollup server is not correct.

Action

Make sure that the user name and password are correct for the device that is being added. Reenter the user name or password and click Add.

BPCUI0510E You entered an invalid time range. The start date and time must be before the end date and time.

Explanation

The start date and time of the time range that you specify must be earlier than the end date and time.

Action

Enter a valid time range, and try the operation again.

BPCUI0511E The following alert name(s) are not unique: *names*.

Explanation

The names of custom alerts must be unique across all resources.

Action

Specify unique names for your custom alerts and try again.

BPCUI0512E Custom alerts already exist for other resources with the following alert name(s): *names*.

Explanation

The names of custom alerts must be unique across all resources.

Action

Specify unique names for your custom alerts and try again.

BPCUI0513E Unable to connect from rollup server *rollup_server* to the repository database.

Explanation

This error typically indicates that the repository database has been shut down. The rollup server must connect to the repository database in order to add or modify subordinate servers.

Action

Please ensure that specified repository RDBMS is up and running.

BPCUI0514E The specified secondary server *secondary_server* is the primary server.

Explanation

This problem might occur when the host name or IP address of the primary server was entered instead of the secondary server.

Action

Enter the correct host name or IP address for the secondary server that you want to add and try again.

BPCUI0515E The duration of the automated probe run window must be at least *minimum_hours* hours.

Explanation

The automated probe window cannot be less than the value that is specified in the message.

Action

Enter the start time and end time so that the duration of the window is at least the amount of time that is specified in the message.

BPCUI0516W The selected subgroups cannot be removed from the general group because they cannot be moved up a level in the groups hierarchy due to name conflicts with the general groups at the higher level.

Explanation

When you remove a subgroup from its parent general group, the subgroup is moved up to the same level in the hierarchy as the parent group. If the subgroup has the same name as an existing general group at the higher level, the subgroup cannot be moved up. The subgroup is not removed from the general group.

Action

Rename the subgroup and try the action again.

BPCUI0519E Authorization has failed because the private key is not valid for the user name that you have specified.

Explanation

The private key and user name that you have provided do not match what has been defined on the cluster.

Action

Ensure that the private key is valid for the specified user name, and log in again.

BPCUI0520E The IP address *ip_address* for the FlashSystem storage system is not the management IP address.

Explanation

The IP address for the FlashSystem storage system that you added is not the management IP address. If you do not use the management IP address, the storage system might be managed incorrectly.

Action

Add the FlashSystem storage system again using the management IP address.

BPCUI0521E The configuration for the report can't be saved.

Explanation

The save operation wasn't completed.

Action

Contact IBM Software Support.

BPCUI0522E Failed to delete a report configuration.

Explanation

An attempt to delete a report configuration failed.

Action

Verify that the database service is up and running.

BPCUI0523E Alerts cannot be defined for this storage system.

Explanation

This error might occur when a storage system is not fully supported for monitoring and data for generating alerts is not being collected.

Action

To verify if the storage system is supported for monitoring, check the support matrix at <http://www.ibm.com/support/docview.wss?uid=swg27047049>.

BPCUI0524E The changes to the report configuration can't be saved.

Explanation

The configuration for the report can't be updated because it couldn't be retrieved from the database.

Action

A new report must be created and configured.

BPCUI0525E The configuration for the report can't be saved because the report title isn't unique.

Explanation

The title of the report must be unique.

Action

Enter a unique title for the report.

BPCUI0527E The action cannot be completed because of an invalid request.

Explanation

The issue should not be seen in the field.

Action

If the problem persists, contact IBM Software Support.

Related reference

-  [Getting support](#)

BPCUI0528E The action cannot be completed because of an invalid file upload request.

Explanation

The issue should not be seen in the field.

Action

If the problem persists, contact IBM Software Support.

Related reference

-  [Getting support](#)

BPCUI0526I The connection test to data source *data source* was successful. A probe is running.

Explanation

The previous state of the monitored device was either unreachable or unknown.

Action

No action is required.

BPCUI0529I The data source *data_Source_Address* is already being managed as a data source for monitoring. No new resources were detected.

Explanation

Only resources that are newly discovered or are not included in a data collection schedule can be added for monitoring.

Action

Enter the IP address or host for a different data source or resource to continue.

If you want to modify resources that are already monitored, go to the list page for the resource. For example, to modify a storage system, in the navigation pane select Storage Resources > Storage Systems. Then, right-click the resource and select View Properties. In the properties notebook, modify the values for the resource.

BPCUI0530I The data source *data_Source_Address* is already being managed as a data source for monitoring. The following new resources were detected:

Explanation

You cannot configure resources that are already being monitored. Only resources that are newly discovered or are not included in a data collection schedule can be added for monitoring.

Action

No action is required.

BPCUI0531E The action cannot be completed because LDAP registry file failed to upload.

Explanation

The LDAP registry file failed to upload.

Action

No action is required.

BPCUI0532E The action failed because of a missing resource.

Explanation

An error occurred while processing a user request.

Action

Ensure that the database repository is running properly. Ensure that all required servers and services are running. Verify that the local area network is available and a firewall is not preventing network access to product services and agents. Try the action again.

If the problem persists, check the log files for error messages that might help determine the problem. For information about the location of log files, check the IBM Knowledge Center.

If none of these actions help resolve the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCUI0533E The LDAP configuration test failed.

Explanation

An error occurred while testing the LDAP configuration.

Action

Ensure that the user ID and password you are using is correct.

BPCUI0534E There was an error executing the collect log process. If this problem persists, you can try collecting and uploading the service logs manually. [Learn More](#).

Explanation

There was a problem while executing the collect log process.

Either the process could not be started, or it has failed unexpectedly.

Action

Collect logs manually: [Collecting Service Logs](#).

After the log package has been created, manually upload the service logs.

BPCUI0535E An FTP connection can not be established. If your organization requires the use of a proxy server, consult the following documentation: [Troubleshooting FTP Transfers](#).

Explanation

A proxy configuration file ibmsdduu.config has been created, but the content is not valid.

Action

Please check and fix the configuration file, and try again.

BPCUI0536E The support data collection failed due to an invalid PMR number format.

Explanation

The used PMR number was not in format nnnnn,nnn,nnn.

Action

Please check the PMR number and try again.

BPCUI0537E The support package could not be created because file system permissions prevent the creation of temporary files.

Explanation

The collection of the support information failed because of problems reading/writing to the file system.

Action

Please check the file system and try again.

BPCUI0538E The support data collection completed creating a support package, but the package could not be uploaded to IBM.

Explanation

The upload of the support package via FTP failed. This might be caused by firewalls preventing the Spectrum Control server to access the internet.

Action

Please do a manual upload of the created service package.

To enable the automatic upload of the support packages in the future, you might have to configure Spectrum Control to use a proxy server. Please see the documentation for details.

BPCUI0539E The support data collection failed with an internal error

Explanation

The creation of the support data package failed due to an internal error.

Action

Try again or try using the CLI command to collect the service information. Please see the documentation for details.

BPCUI0540E The support data collection failed due to an invalid email address format.

Explanation

The used email address was not according to the syntax rules of RFC 822.

Action

Please check the email address and try again.

BPCUI0541E The specified SMI agent was not found. Make sure that the protocol, SMI agent host name or IP address, and port are specified correctly and that the SMI agent is properly configured at that location.

Explanation

The SMI agent specified was not found.

BPCUI0542E A connection was not established. Make sure that the protocol, SMI agent host name or IP address, and port are specified correctly.

Explanation

The attempt to establish a connection failed.

Related reference

- [🔗 Resolving communication issues with Brocade Network Advisor](#)

BPCUI0543E The authentication to the SMI agent failed.

Explanation

The credentials that you supplied for the connection are incorrect.

BPCUI0544E There is a pending delete in process for this SMI agent.

Explanation

The specified SMI agent is currently being deleted.

BPCUI0545E The SMI agent service is not available.

Explanation

The service accessed by the SMI agent is not available.

Action

Make sure that the service accessed by the SMI agent is started and operational.

BPCUI0546E The action cannot be completed because the LDAP registry file could not be updated.

Explanation

The ldapregistry.properties file failed to update.

Action

Check the file ldapregistry.properties file for possible conditions that caused the error. For example, the LDAP registry file is locked or is being used by an application. Correct the error condition and perform the action again.

BPCUI0547E Connection failed. The server might be down or unreachable due to network problems.

Explanation

The test to connect to the server failed.

Action

Test the connection again when the network connection to the server is restored.

BPCUI0548E The add SSL certificate action failed.

Explanation

The script that was used to add the SSL certificate to the webServer keystore file failed.

Action

Verify that you have the correct SSL certificate and try to add the SSL certificate again.

BPCUI0549E The add SSL certificate action failed because of a wrong password.

Explanation

A wrong password was used in the script to add the SSL certificate to the webServer keystore file.

Action

Verify that you have the correct password for the webServer keystore file and try to add the SSL certificate again.

BPCUI0550E The specified storage resource is not valid for the REST API service request.

Explanation

The service request does not support the specified storage resource. For example, you cannot specify a virtual server if you use the Consumption REST API.

Action

Enter a storage resource that is valid for the service request.

BPCUI0551E The file cannot be used because it is not a valid SSL certificate. Select a valid certificate file and try again.

Explanation

The file must be a valid SSL certificate to configure LDAP authentication.

Action

Select a valid SSL certificate file and try again.

BPCUI0554E The SSL certificate download process failed.

Explanation

The process for downloading the SSL certificate from the LDAP server failed.

Action

Try downloading the SSL certificate again. If the problem persists, contact your LDAP administrator.

BPCUI0555E The test connection to the LDAP server failed. Verify that your XML file contains the correct syntax and values and that the LDAP server is running.

Explanation

The test connection to the LDAP server failed.

Action

Verify the settings in the configuration file and test your connection again.

BPCUI0556E An unexpected error occurred creating or updating a support ticket.

Explanation

While executing the action, an error occurred when reading the request data.

BPCUI0557E An invalid request was made when creating or updating a support ticket.

Explanation

While executing the action, an error occurred when reading the request data.

BPCUI0558E This tier name is already in use. Enter a different name .

Explanation

Tier names must be unique so that each tier can be identified.

Action

Enter a unique name for each tier.

BPCUI0559E The custom dashboard was removed by another user. Cancel the action and refresh the page manually.

Explanation

Before the page could be refreshed, the custom dashboard was removed by another user.

Action

Cancel the action or close the dialog, and then refresh the page manually.

BPCUI0600W Can't save the scheduling information for the report because the Data server is offline.

Explanation

The report was saved with the configuration changes that were made, but without the scheduling information.

Complete these actions:

1. From the Home menu, click System Management.
2. In the Components section, click Component Servers.
3. Click Start Server next to the Data server.

When the Data server starts, complete these actions:

1. From the Reports menu, click Reports.
2. In the Custom Report section, click the report, and then click Edit.
3. Click Next, configure the schedule, and save your changes.

Action

Wait a few minutes and then complete these actions:

BPCUI0601I The resource does not have a connection configured. To add a connection to the resource, click Add Switch or Add Fabric.

Explanation

The data source for the resource might have been removed. You must add a connection to the resource again.

Action

Use the Add dialog for the resource to add the connection again.

BPCUI0602E The `osAuthentication` script does not start. The script reported the following error: `script_error`.

Explanation

Action

Fix the reason for failure and try the script again.

BPCUI0603E The connection test to data source `data source` was not successful.

Explanation

The connection to the data source (SMI-S provider) was established but the resource was not found. This problem might occur if the resource was removed and is no longer managed by the SMI-S provider.

Action

In your storage environment, assign the resource to be managed by the SMI-S provider. In IBM Spectrum Control, test the connection again.

BPCUI0604E Can't stop data collection for `entity name`.

Explanation

An error was encountered when stopping the data collection.

Action

Wait a few minutes and try again. If you still can't complete the action, go to your Products and services page (<https://myibm.ibm.com/products-services/>) on IBM Marketplace. Click the down-arrow for the Storage Insights offering, click Support, and then choose an option.

Related reference

-  [Getting support](#)

BPCUI0605E Can't restart data collection for `entity name`.

Explanation

An error was encountered when restarting the data collection.

Action

Wait a few minutes and try again. If you still can't complete the action, go to your Products and services page (<https://myibm.ibm.com/products-services/>) on IBM Marketplace. Click the down-arrow for the Storage Insights offering, click Support, and then choose an option.

Related reference

-  [Getting support](#)

BPCUI0606E The action cannot be completed because there was a failure to create or write into the pending configuration file.

Explanation

An unknown error occurred when trying to create or write the .PendingConfiguration file.

Action

Check that files can be created on the filesystem.

BPCUI0607E The action cannot be completed because there was a failure to read the pending LDAP registry file.

Explanation

An unknown error occurred when trying to read from the ldapregistry.pend file.

Action

Check that the ldapregistry.pend file exists and is readable.

BPCUI0608E The action cannot be completed because there was a failure to get the list of LDAP groups.

Explanation

There was a failure to get LDAP groups. Verify the settings in the uploaded LDAP registry file.

Action

Check the settings in the uploaded LDAP registry file.

BPCUI0609E The Local OS authentication configuration test failed.

Explanation

An error occurred while testing the Local OS authentication configuration.

Action

Ensure that the user ID and password are correct, and the user is member of the group selected.

BPCUI0610E Failed to update modified username for IBMId unique ID.

Explanation

An error occurred trying to update a modified IBMId username.

Action

Try the action again. If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCUI0611E Failed to update IBMid unique ID for IBMid username.

Explanation

An error occurred trying to update a modified IBMid unique ID.

Action

Try the action again. If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

BPCUI0612E The action was not performed due to invalid device credentials for *entity name*.

Explanation

An authentication error occurred trying to execute the selected action for a device. The credentials might need to be updated before the action is retried.

Action

Update the device credentials and test the connection. Try the action again. If the problem persists, contact IBM Software Support.

BPCUI0613E A switch with this host name or IP address is already being monitored.

Explanation

You can add a switch only if it is not already being monitored

Action

Enter a host name or IP address of a switch that is not already being monitored.

BTACD - Database verifier messages for SAN database service

- [BTACD0010I IBM Spectrum Control \(Configuration Service\) initialized successfully.](#)
- [BTACD0011E An error occurred while reading properties from file file name](#)
- [BTACD0012I Properties were successfully read from file file name](#)
- [BTACD0013E Exception occurred while saving the properties file file name.](#)
- [BTACD0014E An error was received while attempting to remove a callback key.](#)
- [BTACD0015E An error was received while attempting to get the database driver driver name.](#)
- [BTACD0016E An error was received while decrypting the database password.](#)
- [BTACD0017E There was an error in setting the WAS Admin password.](#)

BTACD0010I IBM Spectrum Control (Configuration Service) initialized successfully.

Explanation

Configuration Service initialized successfully.

BTACD0011E An error occurred while reading properties from file *file name*

Explanation

There was an error reading from the specified file. This file contains configuration data and is required for the product to initialize.

Action

Verify that the specified file exists and is not corrupt.

- If the file does not exist, restore from a backup copy or reinstall the product.
- If the file exists, verify that it does not have unreadable characters. The file should contain only standard English text characters, integers, and local machine names.

If the problem continues, contact IBM customer support.

Related reference

-  [Getting support](#)

BTACD0012I Properties were successfully read from file *file name*

Explanation

Startup properties were successfully loaded.

BTACD0013E Exception occurred while saving the properties file *file name*.

Explanation

There was an error saving the specified properties file. Updates have not been saved. There might be a problem with the file or you might not have permission to write to the file.

Action

Try to save the properties file again.

If the problem continues, make sure that you have write permission to the file.

BTACD0014E An error was received while attempting to remove a callback key.

Explanation

There was an error removing a callback. Either it has already been deregistered or the key was invalid. This is an internal error that should not affect the product function.

BTACD0015E An error was received while attempting to get the database driver *driver name*.

Explanation

There was an error loading the database driver. Either the driver does not exist or the driver name is incorrect.

Action

Verify that tsnmdbparms.properties lists the driver name as com.ibm.db2.jcc.DB2Driver. Restart the product. If the problem continues, contact IBM customer support.

Related reference

-  [Getting support](#)

BTACD0016E An error was received while decrypting the database password.

Explanation

There was an error decrypting the database password. Either the encrypted password does not exist, it has been corrupted, or the installation has been corrupted.

- If the `tivoli.sanmgmt.jdbc.dbPassword` field in the `tsnmdbparms.properties` file does not contain a 72 character value, then the file is either incorrect or corrupt.
- If the properties file is correct, then the product might not have installed correctly.

Action

Restart the product. If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTACD0017E There was an error in setting the WAS Admin password.

Explanation

The `ChangeWASAdminPass` script file failed.

BTACE - SAN event services messages

- [BTACE0030I IBM Spectrum Control \(Event Service\) initialized successfully.](#)
- [BTACE0031I A SAN Event message was received from: publisher name.](#)
- [BTACE0032E Missing SNMP destination address. SNMP trap was not sent.](#)
- [BTACE0033E An error occurred when trying to send an event to the Tivoli Enterprise Console server.](#)
- [BTACE0034E The Tivoli Enterprise Console server location has not been specified.](#)
- [BTACE0035E JMSEException: exception](#)
- [BTACE0037E An error occurred during the process of forwarding an SNMP trap or sending a Tivoli Enterprise Console event.](#)
- [BTACE0039E Unable to save the SnmpAddress entry to the properties file properties file.](#)
- [BTACE0040E Unable to remove the SnmpAddress entry from the properties file properties file.](#)
- [BTACE0041E Unable to save the SAN Domain ID to the properties file properties file.](#)
- [BTACE0042E Unable to read the SNMP and Tivoli Enterprise Console server address entries from the properties file properties file.](#)
- [BTACE0507E Messaging Service is not running and cannot be used to subscribe or publish events.](#)
- [BTACE0508E A proxy to the Messaging Service could not be obtained.](#)
- [BTACE0509E The proxy to the Messaging Service might not be valid. Cannot publish or subscribe to events.](#)
- [BTACE0510E Unable to create topic topic name.](#)
- [BTACE0511E Unable to create a publisher for topic topic name.](#)
- [BTACE0512E Unable to create a subscriber for topic topic name.](#)
- [BTACE0513E Unable to create a message.](#)
- [BTACE0514E Service service name failed to subscribe to topic topic name.](#)
- [BTACE0515E Publish failed for topic topic name.](#)

BTACE0030I IBM Spectrum Control (Event Service) initialized successfully.

Explanation

SAN Event Services has initialized successfully.

BTACE0031I A SAN Event message was received from: publisher name.

Explanation

SAN Event Service has received a message from the publisher indicated.

BTACE0032E Missing SNMP destination address. SNMP trap was not sent.

Explanation

There is no address to forward SNMP traps to. You can specify an SNMP server to forward events, or traps, to. The SNMP server can be configured to perform actions based on the events it receives.

Action

Specify the address of the SNMP server you would like to forward these events to.

BTACE0033E An error occurred when trying to send an event to the Tivoli Enterprise Console server.

Explanation

An exception occurred when trying to send an event to the Tivoli Enterprise Console server.

Action

Verify that the address of the Tivoli Enterprise Console server is valid.

BTACE0034E The Tivoli Enterprise Console server location has not been specified.

Explanation

There is no location set up for sending the Tivoli Enterprise Console event to.

Action

Specify the address of the Tivoli Enterprise Console server you would like to send the event to.

BTACE0035E JMSException: *exception*

Explanation

An exception was received from the messaging service.

Action

Refer to the logged exception and identify the problem. If the problem continues, contact IBM customer support.

BTACE0037E An error occurred during the process of forwarding an SNMP trap or sending a Tivoli Enterprise Console event.

Explanation

SAN Event Service has received an error while attempting to process an SNMP trap or a Tivoli Enterprise Console event.

Action

Refer to the logged exception and identify the problem. If the problem continues, contact IBM customer support.

BTACE0039E Unable to save the SnmpAddress entry to the properties file *properties file*.

Explanation

There was an error saving the SnmpAddress entry to the properties file. Updates have not been saved. There might be a problem with the file or you might not have permission to write to the file.

Action

Try to save the SnmpAddress entry again and make sure that you have write permission to the properties file *properties file*. If the problem continues, contact IBM customer support.

BTACE0040E Unable to remove the SnmpAddress entry from the properties file *properties file*.

Explanation

There was an error removing the SnmpAddress entry from the properties file. Updates have not been saved. There might be a problem with the file or you might not have permission to write to the file.

Action

Try to remove the SnmpAddress entry again and make sure that you have write permission to the properties file *properties file*. If the problem continues, contact IBM customer support.

BTACE0041E Unable to save the SAN Domain ID to the properties file *properties file*.

Explanation

There was an error saving the SAN Domain ID to the properties file. Updates have not been saved. There might be a problem with the file or you might not have permission to write to the file.

Action

Try to save the SAN Domain ID again and make sure that you have write permission to the properties file *properties file*. If the problem continues, contact IBM customer support.

BTACE0042E Unable to read the SNMP and Tivoli Enterprise Console server address entries from the properties file *properties file*.

Explanation

There was an error reading the SNMP and Tivoli Enterprise Console server address entries from the properties file. There might be a problem with the file or you might not have permission to read the file. This error can also occur if the ConfigService is down.

Action

Make sure the ConfigService is up and that you have permission to read the properties file *properties file*. Then restart the service. If the problem continues, contact IBM customer support.

BTACE0507E Messaging Service is not running and cannot be used to subscribe or publish events.

Explanation

The Messaging Service, which is used by all other services to subscribe to and publish events, is not running. Services cannot subscribe to or publish events.

Action

Restart the product. If the problem persists, contact IBM customer support.

BTACE0508E A proxy to the Messaging Service could not be obtained.

Explanation

A proxy to the Messaging Service could not be obtained.

Action

Review the message log to determine why the proxy could not be obtained. If this problem persists, enable ServiceManager tracing to assist in determining why the proxy could not be obtained.

BTACE0509E The proxy to the Messaging Service might not be valid. Cannot publish or subscribe to events.

Explanation

The proxy to the Messaging Service might not be valid. The Messaging Service might not be running.

Action

Restart the product. If the problem persists, contact IBM customer support.

BTACE0510E Unable to create topic *topic name*.

Explanation

makeTopic tries to construct the full class name of the topic by prefixing the superclass name recursively until the root of the class hierarchy is reached. An error occurred in makeTopic which prevented the full topic name to be constructed.

Action

Enable the messaging middleware trace (san.tesTrcLogger) and view the trace logs to determine why makeTopic failed.

BTACE0511E Unable to create a publisher for topic *topic name*.

Explanation

An error occurred while creating a publisher for the specified topic.

Action

Enable the messaging middleware trace (san.tesTrcLogger) and view the trace logs to determine why the createPublisher method failed.

BTACE0512E Unable to create a subscriber for topic *topic name*.

Explanation

An error occurred while creating a subscriber for the specified topic.

Action

Enable the messaging middleware trace (san.tesTrcLogger) and view the trace logs to determine why the createSubscriber method failed.

BTACE0513E Unable to create a message.

Explanation

An error occurred while creating a message.

Action

Enable the messaging middleware trace (san.tesTrcLogger) and view the trace logs to determine why the createMessage method failed.

BTACE0514E Service *service name* failed to subscribe to topic *topic name*.

Explanation

An error occurred when the service tried to subscribe to the topic.

Action

Enable the messaging middleware trace (san.tesTrcLogger) and view the trace logs to determine why subscribe failed.

BTACE0515E Publish failed for topic *topic name*.

Explanation

An error occurred when the topic was being published.

Action

Enable the messaging middleware trace (san.tesTrcLogger) and view the trace logs to determine why publish failed.

BTACS - Service manager messages

- [BTACS0001I IBM Spectrum Control command line interface initialized successfully.](#)
- [BTACS0002E The -url requires a host:port argument.](#)
- [BTACS0003E Required parameters are not present.](#)
- [BTACS0002I The command line interface is binding to the service.](#)
- [BTACS0003I The bind was successful.](#)
- [BTACS0004E An exception occurred while invoking the service: method name.](#)
- [BTACS0015E Caught exception: value.](#)
- [BTACS0005E Could not deserialize exception: value.](#)
- [BTACS0006E Fault detail: value.](#)
- [BTACS0014E An undeclared exception was encountered: value.](#)
- [BTACS0007E Unrecognized command for service: service name.](#)
- [BTACS0005I Deployed service service name: class=value, scope=value, autostart=value, static=value, order=value.](#)
- [BTACS0006I Undeploying service: service name](#)
- [BTACS0007I Undeployed service: service name](#)
- [BTACS0008E Error undeploying service value : value](#)
- [BTACS0009E Error starting the service name service.](#)
- [BTACS0004I Started service service name.](#)
- [BTACS0011E Service service name was not deployed.](#)
- [BTACS0008I Starting service service name \(timeout number seconds\)](#)
- [BTACS0013E Service service name did not start, interrupting the startup thread.](#)
- [BTACS0009I The service service name was stopped.](#)
- [BTACS0010E An error was encountered while stopping service service name.](#)
- [BTACS0012E Exception was received while stopping service value: value.](#)
- [BTACS0016E Service service name is not available.](#)
- [BTACS0010I Stopping service service name \(timeout number seconds\).](#)
- [BTACS0017E Service service name did not stop, interrupting the startup thread.](#)
- [BTACS0018E A problem was encountered while getting class definition: class definition name.](#)
- [BTACS0019E An interface value in service service name is being ignored.](#)
- [BTACS0020E Error starting service value: value.](#)
- [BTACS0021E Exception getting status from service value: value.](#)
- [BTACS0011I Interrupting monitor thread and waiting for it to exit.](#)
- [BTACS0012I Service Manager shutting down.](#)

- [BTACS0013I Monitoring services \(monitor interval is number seconds\).](#)
- [BTACS0014I Service service name has value.](#)
- [BTACS0015I The Service Manager monitor process is exiting.](#)
- [BTACS0016I There are no services to autostart.](#)
- [BTACS0017I All autostart services have started.](#)
- [BTACS0018I All services are shutting down.](#)
- [BTACS0019I All services have been shut down](#)
- [BTACS0020I Deleted file file name](#)
- [BTACS0021I Unable to delete file file name](#)
- [BTACS0022I Starting autostart services.](#)
- [BTACS0023I An error occurred while starting the service name service.](#)
- [BTACS0024I The properties from file file name were successfully read.](#)
- [BTACS0025E An error occurred while reading properties from file file name.](#)
- [BTACS0026E Login failed: Unknown user name or bad password.](#)
- [BTACS0027E Command failed: Failed to connect.](#)
- [BTACS0028E Command failed.](#)
- [BTACS0029E Failed to connect to <hostname>.](#)
- [BTACS0030E Failed to authenticate host <hostname>.](#)
- [BTACS0031I The server is not registered with the Agent Manager.](#)
- [BTACS0032I Registering with the Agent Manager at host name:port.](#)
- [BTACS0033I The server is renewing credentials with the Agent Manager at host name:port.](#)
- [BTACS0034I The server credentials are current. Agent Manager at host name:port.](#)
- [BTACS0035E The server failed to register with the Agent Manager at host name:port.](#)
- [BTACS0036W The server failed to register with the Agent Manager. It will retry in delay seconds.](#)
- [BTACS0037I The server successfully registered with the Agent Manager.](#)
- [BTACS0038I The server successfully renewed the credentials with the Agent Manager.](#)
- [BTACS0039W The server failed to renew the credentials with the Agent Manager.](#)
- [BTACS0040E This command requires additional arguments.](#)
- [BTACS0041E The command line is not available for service: service name.](#)
- [BTACS0042E Invalid command: CLI command](#)
- [BTACS0043E Failed to authenticate with host hostUrl. Invalid host authentication password.](#)
- [BTACS0044E The server failed to register with the Agent Manager: Incorrect agent registration password.](#)
- [BTACS0045I SERVICE MANAGER COMMANDS](#)
- [BTACS0046I Returns the status of the services.](#)
- [BTACS0047I Service.functionName performed by user at location. Input parameters: input parameters, output parameters: output parameters](#)
- [BTACS0048W Unauthorized request by user at location to perform service.functionName.](#)
- [BTACS0049W Not licensed to perform service.functionName request by user at location.](#)
- [BTACS0050I Waiting for Common Agent services.](#)
- [BTACS0051I The Common Agent services have started successfully.](#)
- [BTACS0052W Failed to create the Common Agent service filter. The Fabric agent will start without waiting for the required Common Agent services.](#)
- [BTACS0053I Agent startup is already in progress.](#)
- [BTACS0054I Invalid Server ID has been provided to update Server Job status.](#)
- [BTACS0055I Agent Manager Registration is set to NO. The server will not register with an AgentManager.](#)
- [BTACS0056I Agent Manager Registration is set to YES. The server will register with the AgentManager.](#)
- [BTACS0057W Error in configuration parameter AgentManager.Registration - default value will be used.](#)
- [BTACS0000I Starting Control Process: value, Device Server RUN ID=value, Job ID=value.](#)

BTACS0001I IBM Spectrum Control command line interface initialized successfully.

Explanation

IBM Spectrum Control command line interface has initialized successfully.

BTACS0002E The -url requires a host:port argument.

Explanation

The -url argument was not entered correctly. It must be in this format: host:port. For example, tiger.wild.com:2210.

Action

Try the command again using the correct syntax for the -url argument.

BTACS0003E Required parameters are not present.

Explanation

The host:port service method [params] required parameters were not entered.

Action

Retry the command with the correct parameters.

BTACS0002I The command line interface is binding to the service.

Explanation

The IBM Spectrum Control command line interface is binding to the service specified on the command line.

BTACS0003I The bind was successful.

Explanation

IBM Spectrum Control command line interface bound to the service specified on the command line.

BTACS0004E An exception occurred while invoking the service: *method name*.

Explanation

An exception was received by the IBM Spectrum Control command line interface while attempting to start the specified service.

Action

Refer to the exception for more information. Contact IBM customer support if the problem continues.

BTACS0015E Caught exception: *value*.

Explanation

An exception was received by the command line interface.

Action

Refer to the exception for more information. Contact IBM customer support if the problem continues.

Related reference

-  [Getting support](#)

BTACS0005E Could not deserialize exception: *value*.

Explanation

An exception was encountered when Service Manager was deserializing an object.

Action

Refer to the exception for more information. Contact IBM customer support.

Related reference

-  [Getting support](#)

BTACS0006E Fault detail: *value*.

Explanation

An error was encountered when Service Manager was deserializing an object.

Action

If the problem persists, shut down and restart IBM Spectrum Control. If this does not correct the error, the problem is probably in network communications.

BTACS0014E An undeclared exception was encountered: *value*.

Explanation

An undeclared exception was encountered by Service Manager.

Action

Refer to the exception for more information. Contact IBM customer support if the problem continues.

Related reference

- [Getting support](#)
-

BTACS0007E Unrecognized command for service: *service name*.

Explanation

Service Manager encountered an internal problem while managing the specified service.

Action

Shut down and restart the program.

BTACS0005I Deployed service *service name*: class=*value*, scope=*value*, autostart=*value*, static=*value*, order=*value*.

Explanation

Service Manager encountered an internal problem while managing a service.

Action

Shut down and restart the program.

BTACS0006I Undeploying service: *service name*

Explanation

Service Manager is undeploying the specified service.

BTACS0007I Undeployed service: *service name*

Explanation

Service Manager successfully undeployed the specified service.

BTACS0008E Error undeploying service *value* : *value*

Explanation

Service Manager encountered a problem while attempting to undeploy the specified service.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0009E Error starting the *service name* service.

Explanation

Service Manager encountered an error while starting the specified service.

Action

Shut down and restart the program. If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0004I Started service *service name*.

Explanation

Service Manager successfully started the specified service.

BTACS0011E Service *service name* was not deployed.

Explanation

The startup or shutdown method was called on a service that has not been deployed.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0008I Starting service *service name* (timeout *number* seconds)

Explanation

Service Manager started the specified service.

BTACS0013E Service *service name* did not start, interrupting the startup thread.

Explanation

Service Manager timed out while starting the specified service.

Action

Change the service.timeout setting in the setup.properties file and restart the program.

BTACS0009I The service *service name* was stopped.

Explanation

Service Manager stopped the specified service.

BTACS0010E An error was encountered while stopping service *service name*.

Explanation

Service Manager encountered an error while stopping the specified service.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0012E Exception was received while stopping service *value*: *value*.

Explanation

Service Manager received an exception while stopping the specified service.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0016E Service *service name* is not available.

Explanation

The requested service is not currently available.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0010I Stopping service *service name* (timeout *number* seconds).

Explanation

Service Manager is stopping the specified service.

BTACS0017E Service *service name* did not stop, interrupting the startup thread.

Explanation

Service Manager timed out while stopping the specified service.

Action

Change the service.timeout setting in the setup.properties file and restart the program.

BTACS0018E A problem was encountered while getting class definition: *class definition name*.

Explanation

Service Manager encountered a problem while getting interface details from the JVM.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0019E An interface *value* in service *service name* is being ignored.

Explanation

Service Manager is ignoring an interface because it could not get details for it from the JVM.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0020E Error starting service *value*: *value*.

Explanation

Service Manager received an exception while stopping the specified service.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0021E Exception getting status from service *value*: *value*.

Explanation

Service Manager received an exception while getting status for the specified service.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0011I Interrupting monitor thread and waiting for it to exit.

Explanation

Service Manager is shutting down - stopping the monitor thread.

BTACS0012I Service Manager shutting down.

Explanation

Service Manager is shutting down.

BTACS0013I Monitoring services (monitor interval is *number* seconds) .

Explanation

Service Manager is monitoring services.

BTACS0014I Service *service name* has *value*.

Explanation

Service Manager is monitoring services.

BTACS0015I The Service Manager monitor process is exiting.

Explanation

The Service Manager monitor process is exiting.

BTACS0016I There are no services to autostart.

Explanation

No services are configured for Service Manager to start.

You can configure services to start automatically. For more information, see the Planning and Installation Guide.

BTACS0017I All autostart services have started.

Explanation

All services configured to autostart have been started.

BTACS0018I All services are shutting down.

Explanation

Service Manager is shutting down all services.

BTACS0019I All services have been shut down

Explanation

Service Manager has shut down all services.

BTACS0020I Deleted file *file name*

Explanation

Service Manager deleted the deployed services file.

BTACS0021I Unable to delete file *file name*

Explanation

Service Manager was unable to delete the specified file.

BTACS0022I Starting autostart services.

Explanation

Service Manager is starting all services that were configured to start automatically.

BTACS0023I An error occurred while starting the *service name* service.

Explanation

Service Manager encountered an error while starting the specified service.

Action

Shut down and restart the program. If the problem continues, contact IBM customer support.

BTACS0024I The properties from file *file name* were successfully read.

Explanation

Startup properties were successfully loaded.

BTACS0025E An error occurred while reading properties from file *file name*.

Explanation

An error occurred while reading from the specified file. This file contains configuration data and is required for the product to initialize.

Action

Verify that the specified file exists and is not corrupt.

- If the file does not exist, restore from a backup copy or reinstall the product.
- If the file exists, verify that it does not have unreadable characters. The file should contain only standard English text characters, integers and local machine names.

If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0026E Login failed: Unknown user name or bad password.

Explanation

The command line was unable to authenticate the specified user with the host operating system.

Action

Verify the user name specified exists in the host operating system and the password is still valid. If this is a Windows machine, confirm that the user that the Fabric service is running under has Act As Part Of The Operating System authority. If running under WebSphere, make sure that WebSphere security is enabled.

BTACS0027E Command failed: Failed to connect.

Explanation

Failed to connect to application.

Action

Failed to make a network connection to the application. Confirm that the application is running.

BTACS0028E Command failed.

Explanation

The command line command failed.

Action

If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0029E Failed to connect to <hostname>.

Explanation

Service Manager was unable to connect to the remote machine.

Action

If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0030E Failed to authenticate host <hostname>.

Explanation

Service Manager was unable to authenticate the remote client.

Action

Reset the host authentication password on the host. If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0031I The server is not registered with the Agent Manager.

Explanation

The server is required to register with the Agent Manager before communicating with any common agents.

BTACS0032I Registering with the Agent Manager at *host name:port*.

Explanation

The server is starting the registration process.

BTACS0033I The server is renewing credentials with the Agent Manager at *host name:port*.

Explanation

The server is renewing credentials with the agent manager because the current credentials will expire soon.

BTACS0034I The server credentials are current. Agent Manager at *host name:port*.

Explanation

The credentials are current.

BTACS0035E The server failed to register with the Agent Manager at *host name:port*.

Explanation

The server failed to register with the Agent Manager.

Action

Check the status of the Agent Manager.

BTACS0036W The server failed to register with the Agent Manager. It will retry in *delay* seconds.

Explanation

The server failed to register with the Agent Manager, but the maximum number of retry attempts has not been reached.

BTACS0037I The server successfully registered with the Agent Manager.

Explanation

The server successfully registered with the Agent Manager.

BTACS0038I The server successfully renewed the credentials with the Agent Manager.

Explanation

The server successfully renewed the credentials with the Agent Manager.

BTACS0039W The server failed to renew the credentials with the Agent Manager.

Explanation

The server failed to renew the credentials with the Agent Manager.

BTACS0040E This command requires additional arguments.

Explanation

More arguments are required for the command line command in order to process the request.

BTACS0041E The command line is not available for service: *service name*.

Explanation

The specified service does not provide a command line interface.

BTACS0042E Invalid command: *CLI command*

Explanation

The specified service does not provide a command line interface.

BTACS0043E Failed to authenticate with host *hostUrl*. Invalid host authentication password.

Explanation

The host authentication password is incorrect.

Action

Reset the host authentication password. If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0044E The server failed to register with the Agent Manager: Incorrect agent registration password.

Explanation

The agent registration password supplied does not match the truststore password.

Action

Reset the agent registration password. If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTACS0045I SERVICE MANAGER COMMANDS

Explanation

Header for the Service Manager command line interface.

BTACS0046I Returns the status of the services.

Explanation

Description of the get status function provided by the Service Manager Command line interface.

BTACS0047I *Service.functionName* performed by user at location. Input parameters: *input parameters*, output parameters: *output parameters*

Explanation

The specified remote procedure call was authorized.

BTACS0048W Unauthorized request by user at location to perform *service.functionName*.

Explanation

The specified remote procedure was not authorized.

BTACS0049W Not licensed to perform *service.functionName* request by user at *location*.

Explanation

The system is not licensed to perform the specified call.

BTACS0050I Waiting for Common Agent services.

Explanation

The agent startup will be performed when the required services have started.

Action

Make sure the Connector service of the Common Agent has started successfully. The service may fail to start if the credentials required for SSL have expired and the Common Agent has not renewed the credentials with the Agent Manager.

BTACS0051I The Common Agent services have started successfully.

Explanation

The required Common Agent services have started. The Fabric agent will proceed with the startup procedure.

BTACS0052W Failed to create the Common Agent service filter. The Fabric agent will start without waiting for the required Common Agent services.

Explanation

The Fabric agent is unable to check if the required Common Agent services have started.

BTACS0053I Agent startup is already in progress.

Explanation

The Fabric agent has already initiated the startup sequence. The current start request will be ignored.

BTACS0054I Invalid Server ID has been provided to update Server Job status.

Explanation

Invalid Server id was provided to update job status. If there are jobs in "running" state, status of those jobs will not be reset to "failed".

BTACS0055I Agent Manager Registration is set to NO. The server will not register with an AgentManager.

Explanation

IBM Spectrum Control is operating in an environment with "AgentManager.Registration" set to NO. The server will not attempt to register with the Agent Manager.

BTACS0056I Agent Manager Registration is set to YES. The server will register with the AgentManager.

Explanation

IBM Spectrum Control is operating in an environment with "AgentManager.Registration" set to YES. The server will attempt to register with the Agent Manager.

BTACS0057W Error in configuration parameter AgentManager.Registration - default value will be used.

Explanation

There was an error in the value of the Configuration parameter "AgentManager.Registration". Allowed values are "yes" and "no".

BTACS0000I Starting Control Process: *value*, Device Server RUN ID=*value*, Job ID=*value*.

Explanation

The specified control process has been started.

Action

None.

BTACS0001I IBM Spectrum Control command line interface initialized successfully.

Explanation

IBM Spectrum Control command line interface has initialized successfully.

BTACS0002I The command line interface is binding to the service.

Explanation

The IBM Spectrum Control command line interface is binding to the service specified on the command line.

BTACS0003I The bind was successful.

Explanation

IBM Spectrum Control command line interface bound to the service specified on the command line.

BTADS/BTAFM/BTAVM/HWN - Job logging messages

- [**BTADS0000I** Starting Discover Process *value* , with Device Server RUN ID *value* , and Job ID *value* .](#)
- [**BTADS0001I** Discover Process with Device Server run ID *value* and job ID *value* is complete.](#)
- [**BTADS0002I** Starting Child Discover Process *value* with Job ID= *value* .](#)
- [**BTADS0003I** The Child Discover Process with Job ID *value* has completed with Status= *value* and Return Code= *value* .](#)
- [**BTADS0004W** The child discovery request with job ID *job_id* completed with status *status* number and return code *value* .](#)
- [**BTADS0005E** The child discovery request with job ID *job_id* completed with status *status* number and return code *value* .](#)
- [**BTADS0010I** Invoking outband scanner *value* on agent *value* .](#)
- [**BTADS0011I** Outband scanner *value* on agent *value* completed successfully.](#)
- [**BTADS0012E** Outband Scanner *value* on agent *value* failed with return code *value* .](#)

- [BTADS0019E An outband scanner failed to capture the scan data.](#)
- [BTADS0020I Processing value data from agent value.](#)
- [BTADS0021W Warning encountered while parsing Fabric XML for job: RUN ID= value , and Job ID= value . value .](#)
- [BTADS0022E Exception encountered while parsing Fabric XML for job: RUN ID= value , and Job ID= value . value .](#)
- [BTADS0023E Fatal error encountered while parsing Fabric XML for job: RUN ID= value , and Job ID= value . value .](#)
- [BTADS0024E Error encountered processing scanner value data from agent value . value .](#)
- [BTADS0025I Running job to discover SMI-S providers through Service Location Protocol: RUN ID= value , Job ID= value .](#)
- [BTADS0026I Service Location Protocol has found value SMI-S providers.](#)
- [BTADS0027E Error encountered by a Service Location Protocol job: RUN ID= value , and Job ID= value . value .](#)
- [BTADS0028W The Device Server Job with RUN ID=; value , Job ID= value , Discover Request= value has been cancelled since it is long running.](#)
- [BTADS0029I Scanner value data from agent value has not changed since last scan.](#)
- [BTADS0030I Invoking inband Scanner value on agent value.](#)
- [BTADS0031I Inband Scanner value on Agent value completed successfully.](#)
- [BTADS0032E Inband Scanner value failed on agent value with Return Code value .](#)
- [BTADS0033E Error invoking value on host value .](#)
- [BTADS0034E Fatal error encountered while persisting the data for job: RUN ID= value , and Job ID= value . value .](#)
- [BTADS0035E The execution of the job failed with: value .](#)
- [BTADS0036I Found SNMP Target at value .](#)
- [BTADS0037E Found SNMP Target at value but it is not persisted in the database. Will NOT perform discovery of information using the address.](#)
- [BTADS0038I Starting scan of SNMP agents from value to value .](#)
- [BTADS0039I Starting probe of detected agents.](#)
- [BTADS0040I Processing of Scanner value data from Agent value completed successfully.](#)
- [BTADS0041I Discover Process with Device Server RUN ID value and Job ID value completed successfully.](#)
- [BTADS0042E Discover Process with Device Server RUN ID value and Job ID value failed with return code value .](#)
- [BTADS0043I Invoking value scanner value on target value .](#)
- [BTADS0044I value scanner value on target value completed successfully.](#)
- [BTADS0045E value Scanner value on target value failed with return code value .](#)
- [BTADS0046I Processing value data from agent value.](#)
- [BTADS0047W The value parser encountered a warning while parsing XML for job with RUN ID= value , and Job ID value . The return code from the parser job is value .](#)
- [BTADS0048E The value parser encountered an exception while parsing XML from job with RUN ID= value , and Job ID= value .The return code from the parser is value .](#)
- [BTADS0049E The value parser for Device Server job with RUN ID= value , and Job ID= value failed. The return code from the parser is value .](#)
- [BTADS0050I Service Location Protocol has found SMI-S provider, value , at address value .](#)
- [BTADS0051I Service Location Protocol has found SMI-S provider, value , at address value , which requires security information to be configured.](#)
- [BTADS0052W Warning encountered while parsing value data from agent value, value.](#)
- [BTADS0053E Exception encountered while parsing value data from agent value, value.](#)
- [BTADS0054E Fatal error encountered while parsing value data from agent value, value.](#)
- [BTADS0055E Outband Scanner value on agent value encountered the presence of a McData i10k. These devices do not report correctly via SNMP and can only be used with SMI-S provider.](#)
- [BTADS0056E Errors in Topology XML generator.](#)
- [BTADS0057E Errors occurred while resolving InterconnectElement and Port relationship.](#)
- [BTADS0058E Errors in creating an entity.](#)
- [BTADS0059E The outband agent target address IP address is not a Cisco device or is invalid.](#)
- [BTADS0060E Outband Scanner value is not responding.](#)
- [BTADS0062E Encountered SQL error value while persisting some data.](#)
- [BTADS0063E The execution of the PM BSP invocation failed with: value .](#)
- [BTADS0063W The performance data collection for the current device is not enabled.](#)
- [BTADS0064I Starting scan of Storage Subsystems from value to value .](#)
- [BTADS0065I Outband and inband agents for fabric\(s\) specified in probe are value](#)
- [BTADS0066I Could not find scanners for agent value](#)
- [BTADS0067I Agent value is configured for no SAN calls and so no scanners will be invoked for this particular agent](#)
- [BTADS0068I Could not retrieve connection information for agent value. Will not be able to invoke scanners for this particular agent](#)
- [BTADS0069I Added inband scanner job with id value discover request value for agent value.](#)
- [BTADS0070I Agent value has not discovered any fabrics and will not be used during the probe.](#)
- [BTADS0071I Invoked inband Scanner value on agent value .](#)
- [BTADS0072I Successfully received response from agent for job value with request id value .](#)
- [BTADS0073E Received error response from agent for job value with request id value. Return code is value.](#)
- [BTADS0074E IP Scan Discovery was canceled due to a hung socket/thread detected. Partial result of the scan will be persisted.](#)
- [BTADS0075E IP Scan Discovery was canceled due to a hung socket/thread detected.](#)
- [BTADS0076I IP Scan Discovery has started for DS, XIV, and IBM SONAS subsystems.](#)
- [BTADS0077I Scanning value out of value IP addresses.](#)
- [BTADS0078I IP Scan Discovery has started for SVC subsystems.](#)
- [BTADS0079I IP Scan Discovery for DS and XIV was done.](#)
- [BTADS0080I IP Scan Discovery for SVC was done](#)
- [BTADS0081I Inband Scanner value for agent address value is not required for probing switches and will not be used.](#)
- [BTADS0082W A first run of a switch probe failed. Additional agents will be used.](#)
- [BTADS0083I The available agents provide a subset of possible features for the probed switch: value](#)
- [BTADS0084I There are no limitations for probing switch value based on the mix of agents that are configured.](#)
- [BTADS0085W A problem was encountered when agent assignments were being determined for the probe. All available agents will be used to collect information about the switch.](#)
- [BTADS0086I The following storage systems were discovered value](#)
- [BTADS0087I IP Scan Discovery did not find any DS8000, SVC, XIV, and IBM SONAS storage systems in the given IP range.](#)
- [BTADS0088I IP Scan Discovery finished.](#)
- [BTADS0089E The Device server is not registered with agent manager. Scanners cannot be used for agent value, value.](#)
- [BTADS0090E There are no agents currently available to probe switch value.](#)
- [BTADS0091I Inband Scanner value for agent address value is currently not running and will not be used.](#)
- [BTADS0092I Inband Scanner value for agent address value is currently disabled from performing fabric functions and will not be used.](#)
- [BTADS0093I Inband Scanner value for agent address value is currently not reachable and will not be used.](#)
- [BTADS0094W The probe for switch value has some limitations.](#)

- [BTADS0095W For switch value some information will not be collected.](#)
- [BTADS0096I The probe limitation can be overcome by configuring an SMI agent to manage fabric value.](#)
- [BTADS0097I The probe limitation can be overcome by configuring SNMP agents to manage switches in fabric value.](#)
- [BTADS0098I The probe limitation can be overcome by configuring a Storage Resource agent to manage fabric value.](#)
- [BTADS0099W The following WWN is not recognized as belonging to a known vendor: value.](#)
- [BTADS0100W Invalid relationships between switches and fabrics were identified. If possible, these relationships will be fixed automatically for the following switches: value.](#)
- [BTADS0101W The discover process that has the Device server run ID value and job ID value completed with one or more warnings.](#)
- [BTADS0102E The probe with the run ID value completed with errors.](#)
- [BTADS0103E No data source is available to probe switch switch_name.](#)
- [BTADS0104E A timeout occurred while processing the request. Try the request again.](#)
- [BTADS0105E A response from the data collector was not received within the specified time.](#)
- [BTADS0106E The requested action on agent agent_name did not complete because the data collector stopped or is not responding. The request failed with error code error_code.](#)
- [BTADS0107W Outband Scanner outband_scanner_name on agent agent_name failed because of another transaction in progress on the switch.](#)
- [BTADS0108E Outband Scanner outband_scanner_name on agent agent_name failed because unexpected data was returned by the switch. Check the trace file for more details.](#)
- [BTADS0109I Outband Scanner outband_scanner_name on agent agent_name did not collect zoning data.](#)
- [BTADS0110I Outband Scanner outband_scanner_name on agent agent_name did not pass write capabilities check.](#)
- [BTADS0111E The probe was unable to collect some details of the switch.](#)
- [BTADS0112E Error encountered while persisting some data. value](#)
- [BTADS0113E Error encountered while processing a probe job. value](#)
- [BTADS0114E The information cannot be saved to the database repository.](#)
- [BTADS0115E The probe failed when collecting information about the resource. The data collector returned the following error status: value.](#)
- [BTAFM0000I Operation op_name processed successfully.](#)
- [BTAFM0100I Initializing Collection.](#)
- [BTAFM0110I Querying the SMI-S provider.](#)
- [BTAFM0113I Collecting for db_table, current_obj of num_objs.](#)
- [BTAFM0114I Probing data for switch switch_name.](#)
- [BTAFM0115I Probing data for port port_name.](#)
- [BTAFM0150I Storing Information.](#)
- [BTAFM0151I The db_table of current_obj num_objs stored.](#)
- [BTAFM0200I Traversing fabric topology.](#)
- [BTAFM0500I The IBM Spectrum Control Device Server service has started successfully.](#)
- [BTAFM0501I The IBM Spectrum Control Device Server service was shut down successfully.](#)
- [BTAFM0502I The IBM Spectrum Control Device Server service provides methods to collect, report and configure the fabric hardware.](#)
- [BTAFM0505I The delete missing function has started.](#)
- [BTAFM0506I The delete missing method was processed in milliseconds milliseconds.](#)
- [BTAFM0723W No blades were discovered for the slot slot.](#)
- [BTAFM2000W Operation op_name partially processed.](#)
- [BTAFM2501W Unable to shut down Device Server Service smoothly.](#)
- [BTAFM4000E Operation op_name failed.](#)
- [BTAFM4001E An internal error occurred.](#)
- [BTAFM4002E Could not get requested information due to an internal error - errorMessage](#)
- [BTAFM4100E Mandatory parameter parameter_name is missing.](#)
- [BTAFM4101E Invalid parameter parameter_name.](#)
- [BTAFM4103E Entity entity_name was not found.](#)
- [BTAFM4104E Attribute attribute_name was not found.](#)
- [BTAFM4105E Computer computer_name was not found.](#)
- [BTAFM4106E Fabric fabric_name was not found.](#)
- [BTAFM4107E Switch switch_name was not found.](#)
- [BTAFM4108E Port port_name was not found.](#)
- [BTAFM4109E Zone set zoneset_name was not found.](#)
- [BTAFM4110E Zone zone_name was not found.](#)
- [BTAFM4111E Zone alias zone_alias_name was not found.](#)
- [BTAFM4112E Zone member zone_member_name was not found.](#)
- [BTAFM4113E Subsystem subsystem_name was not found.](#)
- [BTAFM4114E Host Bus Adapter HBA_name was not found.](#)
- [BTAFM4115E Node node_name was not found.](#)
- [BTAFM4116E Link from port from_port_name to port to_port_name was not found.](#)
- [BTAFM4117E Hub hub_name was not found.](#)
- [BTAFM4118E Router router_name was not found.](#)
- [BTAFM4119E Bridge bridge_name was not found.](#)
- [BTAFM4120E LUN LUN_name was not found.](#)
- [BTAFM4140E Agent Agent_name was not found.](#)
- [BTAFM4141E Scanner scanner_name on agent agent_name was not found.](#)
- [BTAFM4142W Agent agent_name was ignored because the switch switch_name was probed by agent agent1_name.](#)
- [BTAFM4150E Indexed properties property_name don't match.](#)
- [BTAFM4180E Agent to gather sensor and event data is not available for the switch switch_name.](#)
- [BTAFM4200E Credentials not found.](#)
- [BTAFM4300E The connection to the SMI agent for switch switch_name could not be made.](#)
- [BTAFM4301E The invocation of CIM method method_name failed on SMI-S provider SMI-S provider_name. The return code is return_code.](#)
- [BTAFM4302E The invocation of CIM method method_name failed on SMI-S provider SMI-S provider_name with the following exception text: exception_text.](#)
- [BTAFM4303E Received unexpected values from SMI-S provider SMI-S provider_name.](#)
- [BTAFM4304E SMI agent SMI agent_name can not contact switch switch_name.](#)
- [BTAFM4305E The CIM method method_name is not supported on the switch switch_name.](#)
- [BTAFM4306E Could not create connection to SMI-S provider SMI-S provider_name. Reason: reason.](#)
- [BTAFM4307E The username user_name or password is wrong on SMI-S provider SMI-S provider_name.](#)
- [BTAFM4308I Could not create connection to SMI-S provider SMI-S provider_name. Reason: reason. An alternate SMI-S provider will be used.](#)
- [BTAFM4501E No agent is available to configure the zoning on the fabric with ID fabric_name.](#)

- [BTAFM4502E The fabric with ID fabric_name is currently locked by another client of IBM Spectrum Control.](#)
- [BTAFM4503E A token for fabric fabric_name has expired for client client_name.](#)
- [BTAFM4504E The transaction for fabric fabric_name has expired.](#)
- [BTAFM4505E Another transaction is in progress for fabric fabric_name.](#)
- [BTAFM4506E Zone set zoneset_name already exists.](#)
- [BTAFM4507E Zone zone_name already exists.](#)
- [BTAFM4508E Zone alias zone_alias_name already exists.](#)
- [BTAFM4509E Zone member zone_member_name already exists.](#)
- [BTAFM4510E Another job is in progress for fabric fabric_name.](#)
- [BTAFM4550E The Device Server encountered an error accessing the database.](#)
- [BTAFM4600E Unable to start the Device Server service.](#)
- [BTAFM5000E Step failed after collecting Count of collected entities entities for switch switch where entities exist. Continuing with next step.](#)
- [BTAFM5001E No set of fabrics or switches was defined for this probe.](#)
- [BTAFM5002E The SMI agents SMIURL returned an error or can no longer contact the switches.](#)
- [BTAFM5003E Requests to an SMI agent did not correctly collect a set of switches for fabric fabric identity.](#)
- [BTAFM5004E No switch retrieved from the SMI agent for fabric fabric identity.](#)
- [BTAFM5005E No switch found for fabric fabric identity.](#)
- [BTAFM5006E No switch retrieved from database.](#)
- [BTAFM5007E Failed to get CIM entity for fabric fabric_name.](#)
- [BTAFM5008E Failed to get CIM entity for switch switch_name.](#)
- [BTAFM5009E Failed to enumerate CIM entity Entity class name.](#)
- [BTAFM5010E SMI-S provider is not available.](#)
- [BTAFM5011E Failed to get blade for switch Switch name.](#)
- [BTAFM5012E Failed to get physicalpackage for blade with slot number Blade slot name.](#)
- [BTAFM5013E Blade serial number is NULL.](#)
- [BTAFM5014E Step failed after collecting Count of collected entities entities for fabric fabric where entities exist. Continuing with next step.](#)
- [BTAFM5015E Data source could not be retrieved from the IBM Spectrum Control database for fabric fabric where data source exists.](#)
- [BTAFM5016E The selected data source could not be contacted for fabric fabric where data source exists.](#)
- [BTAFM5017E Failed to get fabric for switch Switch name.](#)
- [BTAFM5018E Failed to get CIM entity for port port_name.](#)
- [BTAFM5019E Failed to get switch for port port_name.](#)
- [BTAFM5020E Failed to get blade for port port_name.](#)
- [BTAFM5021E Failed to get CIM entity for blade blade_name.](#)
- [BTAFM5022E Failed to get switch for blade blade_name.](#)
- [BTAFM5023E Failed to discover Fabric and Switch.](#)
- [BTAFM5024E The data source for switch switch_name was not retrieved from the database repository.](#)
- [BTAFM0600I Count of collected entities blades collected for switch switch where entities exist.](#)
- [BTAFM0601I Starting collection of switch blades for switch switch identifier.](#)
- [BTAFM0602I Collection of switch blades completed. Count of collected entities entities collected in total for switch switch identifier.](#)
- [BTAFM0603I Starting collection of switch fports for switch switch identifier.](#)
- [BTAFM0604I Collection of switch fports completed. count of collected entities entities collected in total for switch switch identifier.](#)
- [BTAFM0605I Start probing switch entities switches.](#)
- [BTAFM0606I Start topology probing for fabric fabric entity.](#)
- [BTAFM0609I Count of entities fports collected for switch switch where entities exist.](#)
- [BTAFM0614I The probe task is to probe topology and zone. The probe algorithm is CIM association.](#)
- [BTAFM0616I The probe policy involves discovering segmented or merged fabrics.](#)
- [BTAFM0617I The probe policy doesn't involve discovering segmented or merged fabrics.](#)
- [BTAFM0618I The probe task is to probe topology. The probe algorithm is CIM association.](#)
- [BTAFM0620I Start zone probing for fabric fabric entity.](#)
- [BTAFM0621I Starting collection of zone set for switch switch entity.](#)
- [BTAFM0622I Starting collection of zone for switch switch entity.](#)
- [BTAFM0623I Starting collection of zone alias for switch switch entity.](#)
- [BTAFM0624I Starting collection of zone member from zone alias for switch switch entity.](#)
- [BTAFM0625I Starting collection of zone member and zone alias from zone for switch switch entity.](#)
- [BTAFM0626I Starting collection of zone member from zone for switch switch entity.](#)
- [BTAFM0627I Starting collection of zone set for fabric fabric entity.](#)
- [BTAFM0628I Count of collected entities zone sets collected.](#)
- [BTAFM0629I Collection of zone set completed. Count of collected entities entities collected in total for fabric fabric entity.](#)
- [BTAFM0630I Starting collection of zone for fabric fabric entity.](#)
- [BTAFM0631I Count of collected entities zones collected.](#)
- [BTAFM0632I Collection of zone completed. Count of collected entities entities collected in total for fabric fabric entity.](#)
- [BTAFM0633I Starting collection of zone alias for fabric fabric entity.](#)
- [BTAFM0634I Count of collected entities zone aliases collected.](#)
- [BTAFM0635I Collection of zone alias completed. Count of collected entities entities collected in total for fabric fabric entity.](#)
- [BTAFM0636I Starting collection of zone member from zone alias for fabric fabric entity.](#)
- [BTAFM0637I Starting collection of zone member and zone alias from zone for fabric fabric entity.](#)
- [BTAFM0638I Starting collection of zone member from zone for fabric fabric entity.](#)
- [BTAFM0639I Collection of zone member completed. Count of collected entities entities collected in total for fabric fabric entity.](#)
- [BTAFM0640I Zone probe will discover both active and inactive zone definitions at selected data source datasource name for zone probe.](#)
- [BTAFM0641I Zone probe will discover only active zone sets at data source datasource name for zone probe.](#)
- [BTAFM0654I The port is not switch port.](#)
- [BTAFM0655I The switch profile doesn't support this switch switch_name. No further process to probe this switch.](#)
- [BTAFM0656I Start enumerating entity of association between fabric and zone set at selected data source Url entity.](#)
- [BTAFM0657I Start enumerating entity of association between fabric and zone at selected data source Url entity.](#)
- [BTAFM0658I Start enumerating entity of association between fabric and zone alias at selected data source Url entity.](#)
- [BTAFM0659I Start enumerating entity of association between switch and zone set at selected data source Url entity.](#)
- [BTAFM0660I Start enumerating entity of association between switch and zone at selected data source Url entity.](#)
- [BTAFM0661I Start enumerating entity of association between switch and zone alias at selected data source Url entity.](#)
- [BTAFM0662I Start enumerating associations between virtual fabric and zoning entities at selected data source Url entity.](#)
- [BTAFM0663I Starting collection of switch control processor blades for switch switch identifier.](#)

- [BTAfM0664I Count of collected entities control processor blades collected for switch switch where entities exist.](#)
- [BTAfM0665I Collection of switch control processor blades completed. Count of collected entities entities collected in total for switch switch identifier.](#)
- [BTAfM0666I Checksums for the active and defined Zone Database could not be updated for fabric entity.](#)
- [BTAfM0667E Job id or request id is missing for a SRA job that is been processed.](#)
- [BTAfM0668E Command and/or job timestamp is missing for job id with request id.](#)
- [BTAfM0669I job id with request id was is not found. Device server may have been restarted after job was created.](#)
- [BTAfM0670E could not retrieve output file for job id with request id.](#)
- [BTAfM0671E Another probe of fabric The Name+Nameformat of the fabric is already in progress.](#)
- [BTAfM0672E Device server is not registered with agent manager. Will not be able to invoke scanner on host.](#)
- [BTAfM0673E There are no agents that are currently available to probe fabric.](#)
- [BTAfM0674W No fabric found for event source that is associated with switch with IP address.](#)
- [BTAfM0675E Unable to start parsing of SRA fabric probe data for SRA job id request id file name.](#)
- [BTAfM0676E Error parsing SRA fabric probe data for SRA job id request id file name.](#)
- [BTAfM0677E Unable to connect to SNMP port \(another application may already be connected and forwarding messages\).](#)
- [BTAfM0678I The Name of the switch switch was removed.](#)
- [BTAfM0679I The The Name+Nameformat of the fabric fabric was removed.](#)
- [BTAfM0680E The Name of the switch switch was not removed because it is not missing.](#)
- [BTAfM0681E The The Name+Nameformat of the fabric fabric was not removed because it is not missing.](#)
- [BTAfM0682E An error occurred while checking for access to the database to save new zoning information for fabric to the database.](#)
- [BTAfM0683E Unable to access the database to save zoning information for fabric. Another job is currently saving new zoning information to the database for the same fabric.](#)
- [BTAfM0684I The job is waiting to access the database to save new zoning information for fabric. Another job is currently saving zoning information to the database for the same fabric.](#)
- [BTAfM0685W Host/IP Address is not a switch.](#)
- [BTAfM0686W Switch is not a supported switch.](#)
- [BTAfM0687W The switch does not respond to SNMP queries.](#)
- [BTAfM0688W Cannot communicate with host or IP address.](#)
- [BTAfM0689W No ports were discovered for the switch.](#)
- [BTAfM0690I Collection of data from trunks is completed. Data was collected from count of collected entities trunks.](#)
- [BTAfM0691I Starting collection of data from trunks for switch switch identifier.](#)
- [BTAfM0692I Count of entities trunks collected for switch switch where entities exist.](#)
- [BTAfM0692E A response from the data collector was not received within the specified time.](#)
- [BTAfM0693E A response from the data collector was not received. The request failed with return code return_code](#)
- [BTAfM0694W Zoning data cannot be collected because there is a transaction in progress on the switch key](#)
- [BTAfM0695E The switch key is returning unexpected data.](#)
- [BTAfM0696E Zone set zoneset_name is already active.](#)
- [BTAfM0697E Zone set zoneset_name is already inactive.](#)
- [BTAfM0698E On the switch switch_name VSAN vsan_name was not found.](#)
- [BTAfM0699E The switch key did not return zoning data.](#)
- [BTAfM0700E Duplicate entries for the same switch: switch.](#)
- [BTAfM0701E Current active full zone configuration is not synchronized with the zone configuration on the switch switch_name for VSAN vsan_name.](#)
- [BTAfM0702E You cannot monitor Brocade Access Gateway switches without Network Advisor.](#)
- [BTAfM0703I Waiting for probes of other Access Gateway switches to complete.](#)
- [BTAfM0704W Distributing zone configuration across all the switches for VSAN vsan_name did not succeed on the switch switch_name.](#)
- [BTAfM0705W Zone data collection after zone changes were made failed on the switch switch_name.](#)
- [BTAfM0706E The fabric probe was unable to collect some details of the blades on the switches.](#)
- [BTAfM0707I You cannot use IBM Spectrum Control to make zoning changes for provisioning on switch switch_name.](#)
- [BTAfM0708E The probe was unable to collect some details of the switches.](#)
- [BTAfM0709I Started to process information for fabric fabric_name.](#)
- [BTAfM0710I Started to process information for switch switch_name.](#)
- [BTAfM0711I Started to process information for discovered switches.](#)
- [BTAfM0712I Started to process information for a switch blade.](#)
- [BTAfM0713I Started to process information for a switch zone set.](#)
- [BTAfM0714I Started to process information for switch ports.](#)
- [BTAfM0715E Error occurred while processing information for fabric fabric_name.](#)
- [BTAfM0716E Error occurred while processing information for virtual fabric virtual_fabric_name.](#)
- [BTAfM0717E Error occurred while processing information for switch switch_name.](#)
- [BTAfM0718E Error occurred while processing information for discovered switches.](#)
- [BTAfM0719E Error occurred while processing information for logical switches.](#)
- [BTAfM0720E Error occurred while processing information for active zone set active_zone_set_name.](#)
- [BTAfM0721E Error occurred while processing information for inactive zone set inactive_zone_set_name.](#)
- [BTAfM0722E Error occurred while processing information for port port_name.](#)
- [BTAQE1107E InbandScanHandler failed to start InbandScanner scanner_name on managed host target.](#)
- [BTAQE1108E InbandScanHandler failed to get callback information for InbandScanner scanner_name on managed host target.](#)
- [BTAQE1112E During an outband scan, the scanner scanner_name was unable to identify the target host target.](#)
- [BTAQE1113E Unable to invoke an Outband scan scanner_name on target target.](#)
- [BTAQE1114E OutbandScannerHandler received invalid callback information for Outband scanner scanner_name on target target.](#)
- [BTAQE1115E The outband scanner scanner_name did not return the SAN ID on target target.](#)
- [BTAfM0001I The operation Name of the operation processed successfully.](#)
- [BTAfM0002I The Web service call Name of the operation processed successfully.](#)
- [BTAfM0003I Data source Name of the datasource successfully added.](#)
- [BTAfM0004I Data source Name of the datasource successfully deleted.](#)
- [BTAfM0005I Data source Name of the datasource successfully modified.](#)
- [BTAfM0006I Discovery on data source Name of the datasource has started.](#)
- [BTAfM0007I Discovery on data source Name of the datasource completed successfully.](#)
- [BTAfM0008I Probe of hypervisor Name of the Hypervisor has started.](#)
- [BTAfM0009I Probe of hypervisor Name of the Hypervisor completed successfully.](#)
- [BTAfM0010I A connection test to data source Name of the data source has started.](#)
- [BTAfM0011I The Connection test to data source Name of the data source completed successfully.](#)
- [BTAfM0012I Hypervisor Name of the Hypervisor discovered/rediscovered.](#)

- [BTAVM0013I Discovery: Hypervisor Name of the hypervisor will not be discovered as it is managed by another data source.](#)
- [BTAVM0014I Discovery: Hypervisor Name of the hypervisor will not be discovered as it itself is registered as a data source.](#)
- [BTAVM0015I Collection of the physical storage configuration for hypervisor Name of the hypervisor has started.](#)
- [BTAVM0016I Collection of the physical storage configuration for hypervisor Name of the hypervisor completed successfully.](#)
- [BTAVM0017I Collection of the logical storage configuration for hypervisor Name of the hypervisor has started.](#)
- [BTAVM0018I Collection of the logical storage configuration for hypervisor Name of the hypervisor completed successfully.](#)
- [BTAVM0019I Collection of the virtual machines configuration for hypervisor Name of the hypervisor has started.](#)
- [BTAVM0020I Collection of the virtual machines configuration for hypervisor Name of the hypervisor completed successfully.](#)
- [BTAVM0021I The probe of name of the hypervisor found number of physical disks physical disks.](#)
- [BTAVM0022I The probe of name of the hypervisor found number of logical volumes logical volumes.](#)
- [BTAVM0023I The probe of name of the hypervisor found number of virtual machines virtual machines.](#)
- [BTAVM0024I The Name of the hypervisor hypervisor was removed.](#)
- [BTAVM0025I VMWare Cluster Name of the Cluster discovered/rediscovered.](#)
- [BTAVM1301I The probe of name of the hypervisor could collect partial information only for the disk with the device name Device name of the disk.](#)
- [BTAVM1302I LUN correlation is not supported for disk with device name Device name of the disk, vendor: Vendor name, model: model name, for hypervisor hypervisor name.](#)
- [BTAVM1503E An internal error occurred: Text describing the internal error.](#)
- [BTAVM2001E The mandatory parameter Name of the mandatory parameter which is missing is missing.](#)
- [BTAVM2002E Invalid parameter Name of the parameter which was invalid.](#)
- [BTAVM2003E A database error was encountered during database query or insert.](#)
- [BTAVM2004E Cannot connect to the database repository.](#)
- [BTAVM2006E The operation Name of the operation that failed failed for the following reason: Reason of the failure.](#)
- [BTAVM2007E The Web service call Name of the operation failed for the following reason: Reason of the failure.](#)
- [BTAVM2008E The product Name of the unsupported product is not supported.](#)
- [BTAVM2010E The user name or password is invalid for Address of the host](#)
- [BTAVM2011E The operation Name of the timed out operation could not complete within the time limit of Timeout threshold in milliseconds milliseconds.](#)
- [BTAVM2012E An error occurred while trying to establish secure communication over SSL.](#)
- [BTAVM2013E The Add Device wizard could not add the Name of the data source data source.](#)
- [BTAVM2014E The deletion of data source Name of the data source failed.](#)
- [BTAVM2015E The modification of data source Name of the data source failed.](#)
- [BTAVM2016E Discovery on data source Name of the datasource failed.](#)
- [BTAVM2017E Probe of the hypervisor Name of the Hypervisor failed.](#)
- [BTAVM2018E IBM Spectrum Control can't connect to the data source Name of the datasource.](#)
- [BTAVM2201E Probe: An error occurred during the collection of the physical storage configuration.](#)
- [BTAVM2202E Probe: An error occurred during the collection of the logical storage configuration.](#)
- [BTAVM2204E Probe: An error occurred during the collection of the virtual machine configuration.](#)
- [BTAVM2206E Discovery: the hypervisor Name of the hypervisor will not be discovered because its version is not supported.](#)
- [BTAVM2207E Calculation of the summary data for the hypervisor Name of the hypervisor failed.](#)
- [BTAVM2208E Unable to obtain the hypervisor version\(s\) from the datasource Name of the datasource.](#)
- [BTAVM2209E Unable to obtain information about other Virtual Centers managing the hypervisor\(s\) of datasource Name of the datasource.](#)
- [BTAVM2210W Error getting LUN definition data for the disk with the device name Device name of the disk, storage subsystem vendor: Vendor name, model: model name, for hypervisor hypervisor name.](#)
- [BTAVM2211E Probe: Virtualization Manager failed to get the VMWare VI data source for the hypervisor Name of the hypervisor from the database.](#)
- [BTAVM2212E Probe: The hypervisor Name of the hypervisor is not available on the VMWare VI datasource Name of the datasource.](#)
- [BTAVM2213E Data source Name of the datasource is disconnected from Virtual Center.](#)
- [BTAVM2214E The probe job encountered an NFS file system while probing ESX server {0}. IBM Spectrum Control currently does not support probes of ESX servers with NFS file systems. The probe job for this ESX server has been stopped. Probes of other ESX servers that are included in this probe job will continue.](#)
- [BTAVM2215W Unsupported storage subsystem disk with device name Device name of the disk, vendor: Vendor name, model: model name, for hypervisor hypervisor name with hypervisor version less than 3.5.0.](#)
- [BTAVM2216E Unable to get keystore instance.](#)
- [BTAVM2217E Unable to load keystore file.](#)
- [BTAVM2218E Unable to set certificate entry in keystore file.](#)
- [BTAVM2219E Unable to open keystore for writing.](#)
- [BTAVM2220E Unable to close keystore file.](#)
- [BTAVM2221E Unable to acquire lock on keystore file.](#)
- [BTAVM2222E Unable to store certificate in keystore file.](#)
- [BTAVM2223E Unable to release lock on keystore file.](#)
- [BTAVM2224E Unable to decrypt keystore password.](#)
- [BTAVM2225E Unable to open keystore for reading.](#)
- [BTAVM2226E Certificate already exists in keystore.](#)
- [BTAVM2227E host_address hypervisor is already being monitored and could not be added.](#)
- [BTAVM2228E Missing host name.](#)
- [BTAVM2229E Missing certificate.](#)
- [BTAVM2230E Cannot create keystore directory.](#)
- [BTAVM2231E Cannot download the certificate from Data Source Name of the data source.](#)
- [BTAVM2232E Cannot connect to the Name of the data source data source.](#)
- [BTAVM2233E Cannot download the certificate from the port.](#)
- [BTAVM2234E The hypervisor name hypervisor was not removed because IBM Spectrum Control is running other actions on the device.](#)
- [BTAVM2235E Unable to obtain the cluster\(s\) from the datasource Name of the datasource.](#)
- [BTAVM2236W Subsequent steps of probe process may not be able to collect data for the hypervisor Name of the hypervisor because the hypervisor is in critical state.](#)
- [BTAVM2237E Datastore Browser Task failed for hypervisor Name of the hypervisor, datastore Name of the datastore with error: Error](#)
- [BTAVM2238E The registration of the vSphere Web Client extension for IBM Spectrum Control has started on Name of the vCenter server.](#)
- [BTAVM2239E The registration of the vSphere Web Client extension for IBM Spectrum Control did not extract the extension package.](#)
- [BTAVM2240E The registration of the vSphere Web Client extension for IBM Spectrum Control did not complete while updating the VASA web archive file, vasa.war, with the IBM Spectrum Control server configuration.](#)
- [BTAVM2241E The registration of the vSphere Web Client extension for IBM Spectrum Control completed.](#)
- [BTAVM2242E Unable to register IBM Spectrum Control as an extension on the vCenter server Name of the vCenter server. The validation of input values did not complete.](#)

- [BTAVM2243E Unable to register IBM Spectrum Control as an extension on the vCenter server Name of the vCenter server. Could not authenticate with the vCenter server.](#)
- [BTAVM2244E The registration of the vSphere Web Client extension for IBM Spectrum Control did not complete.](#)
- [BTAVM2245E Unable to connect to the vCenter Server Name of the datasource.](#)
- [BTAVM2246E Unable to configure the vCenter Server.](#)
- [BTAVM2247E The registration of the vSphere Web Client extension for IBM Spectrum Control did not delete the temporary directory Name of the directory.](#)
- [BTAVM2248E The registration of IBM Spectrum Control as a VASA provider did not complete.](#)
- [BTAVM2249E Automatic registration of IBM Spectrum Control as a VASA provider is not supported for vCenter Server version 5.0 and earlier.](#)
- [BTAVM2250E IBM Spectrum Control is already registered as a VASA provider for vCenter Server server_name. Register IBM Spectrum Control as a VASA provider manually in the vSphere Web Client to update the credentials.](#)
- [BTAVM2251E One or more third-party VASA providers are already registered with the vCenter Server. IBM Spectrum Control VASA provider was not registered. Register IBM Spectrum Control as a VASA provider manually.](#)
- [BTAVM2252E The registration of IBM Spectrum Control as a VASA provider has started on Name of the vCenter server.](#)
- [BTAVM2253E The registration of IBM Spectrum Control as a VASA provider has completed.](#)
- [BTAVM2254E The registration of the vSphere Web Client extension for IBM Spectrum Control did not complete. The current session is invalid.](#)
- [BTAVM2255E The registration of IBM Spectrum Control as a VASA provider did not complete. The current session is invalid.](#)
- [BTAVM2256W Could not determine the host for VM with ID: host id and Name: Vendor name. Check if the same mac address is used on other computers.](#)
- [BTAVM2257I Found number of files files on name of datastore of name of the hypervisor.](#)
- [BTAVM2258I The probe of name of the hypervisor found number of controllers controllers.](#)
- [BTAVM2259I Collecting file system details for hypervisor Name of the hypervisor.](#)
- [BTAVM2260I Collecting list of files for hypervisor Name of the hypervisor.](#)
- [BTAVM2261I Collecting logical volumes for hypervisor Name of the hypervisor.](#)
- [BTAVM2262I Collecting disk partition for hypervisor Name of the hypervisor.](#)
- [BTAVM2263I Files details for Name of the datastore being collected by id of the Hypervisor.](#)
- [BTAVM2264I Files details for Name of the datastore were collected by id of the Hypervisor on timestamp.](#)
- [BTAVM2265E Invalid host name or IP address.](#)
- [BTAVM2266E The connection information cannot be updated because it points to another device.](#)
- [BTAVM2268E The connection information cannot be updated because IBM Spectrum Control cannot determine if the hypervisor is managed by the Name of the data source data source.](#)
- [BTAVM2269E The connection information cannot be updated because a data source with this host name or IP address is already present.](#)
- [BTAVM2270E The connection information cannot be updated because it doesn't point to a data source of the same type \(vCenter/ESX\).](#)
- [BTAVM2271W The hypervisor Name of the Hypervisor cannot be discovered because its connection state is "Connection State".](#)
- [BTAVM2272E The user User Name does not have the privilege to browse the datastore Name of the Datastore.](#)
- [HWN020001I Operation Name of the operation processed successfully.](#)
- [HWN020002E Mandatory parameter Name of the mandatory parameter which is missing missing](#)
- [HWN020003E Invalid parameter Name of the parameter which was invalid](#)
- [HWN020101E The external process terminated unexpectedly.](#)
- [HWN020102W The external process was canceled per users request.](#)
- [HWN020103E The external process exceeded the timeout limit and was canceled.](#)
- [HWN020104E The external process could not be started.](#)
- [HWN020105E The data collector is not responding to the server.](#)
- [HWN020106E An external process was cancelled by the data collector.](#)
- [HWN021503E The action cannot be completed](#)
- [HWN021504E Entity The ID of the entity was not found.](#)
- [HWN021508E Credentials not found](#)
- [HWN021514E The invocation of CIM method Name of method failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method](#)
- [HWN021515E The invocation of CIM method Name of method failed on SMI-S provider Name of SMI-S provider with the following exception text: Exception text](#)
- [HWN021516E The LSS specified LSS name on subsystem Name of subsystem is already at the maximum volume number \(255\). Volume creation can not be done on this LSS. please select a different one.](#)
- [HWN021517E The connection to SMI-S provider for storage system VPD of the storage system could not be made.](#)
- [HWN021520E The attribute Name of the attribute was not found.](#)
- [HWN021522E Host port The WWPN of the host port not assigned to Volume The PK of the volume](#)
- [HWN021524E Indexed Properties Names don't match](#)
- [HWN021529E An SMI-S provider has reported unexpected values: IP and port of SMI-S provider.](#)
- [HWN021530E The Volume - Port mapping can not be created. There are existing mappings that prevent this combination. VolumeCOP: The ID of the volume ,Port: The WWPN of the port that should be mapped to the volume](#)
- [HWN021531E SMI-S provider The IP and port of the SMI-S provider can not reach storage system The VPD of the storage system](#)
- [HWN021535E There is not enough space left in the storage pool The primary key of the Poolon storage system The VPD of the storage system to create a volume of The requested volume size bytes.](#)
- [HWN021536E The CIM method The CIM method that is not supported. is not supported on the storage system The VPD of the storage system](#)
- [HWN021537E Could not create connection to SMI-S provider The IP and port for the SMI-S provider..Reason: The exception returned by the SMI-S provider.](#)
- [HWN021538E The username The username that was used to connect to the SMI-S provider. or password is wrong on SMI-S provider The IP and port for the SMI-S provider.](#)
- [HWN021539E The SVC with IP The IP of the SVC. which is managed by SMI-S provider The IP and port for the SMI-S provider. can not be discovered. The status is The status of the SVC..](#)
- [HWN021540E The invocation of CIM method Name of method failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method. Details provided by the SMI-S provider : Description of Returncode](#)
- [HWN021600W Operation Name of the operation. partially processed.](#)
- [HWN021601E The operation\(s\) Operation_names failed.](#)
- [HWN021602E It is necessary to specify target ports for storage device VPD of the storage subsystem](#)
- [HWN021603W More storage volumes and ports than specified will loose access](#)
- [HWN021604E WWPNs and storage volumes to be unassigned not completely specified. Assigned WWPNs: All WWPNs that are assigned to the volumes in the host port collection. ,missing WWPNs: The WWPNs that are assigned but were not specified in the input parameter in the method unassign . Storage volumes to be unassigned not completely specified. Assigned storage volumes: Lists all storage volumes that are really assigned to the WWPNs. }, missing storage volumes: The storage volumes that are really assigned but were not specified in the input parameter in the method unassign](#)
- [HWN021605I More storage volumes and ports than specified will gain access.](#)
- [HWN021606E WWPNs and storage volumes to be assigned not completely specified. Missing WWPNs: The WWPNs that need to be assigned but were not specified in the input parameter. .Storage volumes to be assigned not completely specified. Missing storage volumes: The storage volumes that need to be assigned but were not specified in the input parameter.](#)

- [HWN021607E The client type the client type with description the client description is not supported on SMI-S provider the SMI-S provider IP and port for storage subsystem the subsystem ID of volumes the volumeIDs of the subsystem which were passed in](#)
- [HWN021608E The target port the target port ID does not belong to storage subsystem the subsystem ID of volumes the volumeIDs of the subsystem which were passed in](#)
- [HWN021609E There is not enough space left in the storage pool The primary key of the Pool on storage system The VPD of the storage system to create The number of volumes to create volumes of The total size needed bytes total.](#)
- [HWN021610E The specified size The size of the volume to create is not supported on pool The storage pool ID Size has to be dividable by Divisor returned by getSupportedSizeRange and in between Minimum returned by getSupportedSizeRange and Maximum returned by getSupportedSizeRange](#)
- [HWN021611E Volume The volume ID has mappings, it can not be deleted.](#)
- [HWN021612E The mapping between volume The volume ID and port The initiator port wwpn exists already](#)
- [HWN021613E The WWPEN The WWPEN not found can not be found on subsystem The subsystem](#)
- [HWN021614E The WWPENs The WWPENs without mappings have no mappings on storage system The storage system](#)
- [HWN021615E WWPENs WWPENs that can not share mappings can not share mappings on storage system Storage system}. There are existing mappings that prevent this.](#)
- [HWN021616E Volumes VolumeIDs can not share mappings on storage system Storage system }. There are existing mappings that prevent this.](#)
- [HWN021617E The stored data for storage system The storage system is not in sync with the environment. Rerun data collection.](#)
- [HWN021618E Modifying target ports is not supported by subsystem the subsystem.](#)
- [HWN021619E Modifying the target ports for mapping of initiator port initiator port WWPEN and volume volume name will also modify the target ports of the following mappings: port - volume list](#)
- [HWN021620I Modifying the target ports for mapping of initiator port initiator port WWPEN and volume volume name will modify the target ports of more mappings than specified.](#)
- [HWN021621E It is not supported to modify the target ports of existing mappings and create new mappings in one step. Modify the existing mappings first and then create the new mappings. Existing mappings: port - volume list](#)
- [HWN021622I Started modification of the assignment of volume VolumeID on subsystem Subsystem to initiator port WWPEN . Target ports to add: target ports to add Target ports to remove: target ports to remove](#)
- [HWN021623I Finished modification of the assignment of volume VolumeID on subsystem Subsystem to initiator port WWPEN . Target ports to add: target ports added Target ports to remove: target ports removed](#)
- [HWN021624E The modification of the assignment of volume VolumeID on subsystem Subsystem to initiator port WWPEN failed. Target ports to add: target ports to add Target ports to remove: target ports to remove](#)
- [HWN021650E A timeout occurred while connecting to SMI-S provider SMI-S provider IP and port.](#)
- [HWN021651E Job on SMI-S provider SMI-S provider IP and Port in format IP:Port failed. Job Status: Job status . Error code is Error code , error description: Error description . Check IBM Spectrum Control and SMI-S provider logs.](#)
- [HWN021652E The process has timed out. Check the IBM Spectrum Control log files for more information.](#)
- [HWN021653E The attribute Name of the attribute was not found.](#)
- [HWN021654E Pool ID was not found.](#)
- [HWN021655E Volume ID The ID of the volume was not found.](#)
- [HWN021656E Port ID The ID of the port was not found.](#)
- [HWN021657E Subsystem ID The ID of the subsystem was not found.](#)
- [HWN021658E Managed Disk ID The ID of the MDisk was not found.](#)
- [HWN021659E SMI-S provider The ID of the SMI-S provider was not found](#)
- [HWN021660E IO Group The SVC IO Group was not found.](#)
- [HWN021661E Extent The storage extent external key was not found.](#)
- [HWN021662E Physical volume The physical volume external key was not found.](#)
- [HWN021670E The client type the client type with description the client description is not unique on SMI-S provider the SMI-S provider IP and port } for storage subsystem the subsystem ID of volumes the volumeIDs of the subsystem which were passed in](#)
- [HWN021671I The storage system The storage system was deleted from the database](#)
- [HWN021672E The storage system name storage system was not removed because other monitoring actions are running on the device.](#)
- [HWN021673E The probe job on SMI-S provider SMI-S provider IP and Port in format IP:Port did not complete within the time limit of Microseconds microseconds. The job is Percent complete percent complete. Check the SMI-S provider log for job status. Job information: JobCOP . Run the probe job again after the current job has completed.](#)
- [HWN021674E Job on SMI-S provider SMI-S provider IP and Port in format IP:Port returned unexpected results. Job information: JobCOP Job status: JobState , status description: JobStatus Check SMI-S provider log. Redo probe if the job completed.](#)
- [HWN021675I Started creation of volume with size Size in pool Pool on subsystem Subsystem](#)
- [HWN021676I Volume creation completed successfully. New volume VolumeID created with size Size in pool Pool on subsystem Subsystem .](#)
- [HWN021677E Volume creation failed. The volume of size Size in pool Pool on subsystem Subsystem could not be created.](#)
- [HWN021678I Started assignment of volume VolumeID on subsystem Subsystem to initiator port WWPEN .](#)
- [HWN021679I Finished assignment of volume VolumeID on subsystem Subsystem to initiator port WWPEN .](#)
- [HWN021680E The assignment of volume VolumeID on subsystem Subsystem to initiator port WWPEN failed.](#)
- [HWN021681I Started unassignment of volume VolumeID on subsystem Subsystem to initiator port WWPEN .](#)
- [HWN021682I Finished unassignment of volume VolumeID on subsystem Subsystem to initiator port WWPEN .](#)
- [HWN021683E The unassignment of volume VolumeID on subsystem Subsystem to initiator port WWPEN failed.](#)
- [HWN021684I Started deletion of volume VolumeID on subsystem Subsystem .](#)
- [HWN021685I Volume deletion completed successfully. Volume VolumeID on subsystem Subsystem was deleted.](#)
- [HWN021686E Volume deletion failed. Volume VolumeID on subsystem Subsystem could not be deleted.](#)
- [HWN021687I Started modification of Pool Pool display name on subsystem Subsystem display name .](#)
- [HWN021688I Pool modification completed successfully. Pool Pool display name on subsystem Subsystem display name was modified.](#)
- [HWN021689E Pool modification failed. Pool Pool display name on subsystem Subsystem display name could not be modified.](#)
- [HWN021690I Started creation of number volumes volumes with size Size in pool Pool on subsystem Subsystem](#)
- [HWN021691I Created number volumes out of total number volumes volumes with size Size in pool Pool on subsystem Subsystem](#)
- [HWN021692E Volume creation failed. Created number volumes out of total number volumes volumes with size Size in pool Pool on subsystem Subsystem](#)
- [HWN021693W Warning: The task succeeded, but the database update failed. Run probe to update the database.](#)
- [HWN021700I Enumerating CIM Associator The CIM association name which is being enumerated. for The name of the DB table which will be populated as result of this query.](#)
- [HWN021701I Enumerating CIM Class The CIM class name which is being enumerated. for The name of the DB table which will be populated as result of this query.](#)
- [HWN021702I Querying SMI-S provider](#)
- [HWN021703I Task starting on SMI-S provider Identifier of the SMI-S provider..](#)
- [HWN021708I Initializing Collection for storage system storage system identification.](#)
- [HWN021709I Collection for storage system storage system identification completed.](#)
- [HWN021710I Discovering devices for SAN Volume Controller The VPD of the SAN Volume Controller.](#)

- [HWN021711I Discovery devices for SAN Volume Controller The VPD of the SAN Volume Controller. failed with error message The exception which has occurred.](#)
- [HWN021712I Collecting Nodes for storage system storage system identification.](#)
- [HWN021713I Collecting fibre channel ports for storage system storage system identification.](#)
- [HWN021714I Collecting volumes for storage system storage system identification.](#)
- [HWN021715I Traversing host to volume assignments for storage system storage system identification.](#)
- [HWN021716I Collecting pools and volumes for storage system storage system identification.](#)
- [HWN021717I Collecting volume settings for storage system storage system identification.](#)
- [HWN021718I Collecting client setting data for storage system storage system identification.](#)
- [HWN021719I Perform collection post process tasks for storage system storage system identification.](#)
- [HWN021720I Flash enclosure is missing drive flash_drive_identifier.](#)
- [HWN021724W SMI-S provider SMI-S provider identifier manages device\(s\) of type device_type which is supported through the native device interface or SNMP only.](#)
- [HWN021725I IBM Spectrum Control discovered/rediscovered a device with name Identifier of the device. on SMI-S provider Identifier of the SMI-S provider..](#)
- [HWN021726I IBM Spectrum Control discovered/rediscovered no device on SMI-S provider Identifier of the SMI-S provider..](#)
- [HWN021727I IBM Spectrum Control discovery starting on SMI-S provider Identifier of the SMI-S provider..](#)
- [HWN021728I IBM Spectrum Control discovery on SMI-S provider Identifier of the SMI-S provider. is complete.](#)
- [HWN021729W IBM Spectrum Control discovery of Device type value is not supported.](#)
- [HWN021730W IBM Spectrum Control discovery of device value with code level value is not supported on SMI-S provider Identifier of the SMI-S provider..](#)
- [HWN021731I Probing Volumes for Storage System: value.](#)
- [HWN021732I Number of Volumes Found Currently: value. Continuing to Probe Volumes.](#)
- [HWN021733I value Volumes Found.](#)
- [HWN021734I Probing Disks for Storage System: value.](#)
- [HWN021735I Number of Disks Found Currently: value. Continuing to Probe Disks.](#)
- [HWN021736I value Disks Found.](#)
- [HWN021737I Probing Virtual Disks for Cluster: value](#)
- [HWN021738I Number of Virtual Disks currently found: value. Continuing to probe Virtual Disks.](#)
- [HWN021739I value Virtual Disks found.](#)
- [HWN021740I Probing Views of Host Initiator access to Volumes.](#)
- [HWN021741I value Views Found.](#)
- [HWN021742E The SMI-S provider SMI-S provider URL is not managing storage subsystems.](#)
- [HWN021743E The SMI-S provider SMI-S provider URL is not managing switches.](#)
- [HWN021744E Cannot connect to a resource because of an SSL certificate error. Troubleshooting information: http://www.ibm.com/support/docview.wss?uid=swg21976237](#)
- [HWN021745I Cannot connect to a resource because of an SSL certificate error. Troubleshooting information: http://www.ibm.com/support/docview.wss?uid=swg21976237. An alternate resource will be used.](#)
- [HWN021746W SMI-S provider Identifier of the SMI-S provider. manages Cisco device types through SNMP only.](#)
- [HWN021747E Unable to add the specified switch by using SNMP. The switch is a Brocade switch and can be added only by using an SMI agent.](#)
- [HWN021800E Failed to get a database connection.](#)
- [HWN021801E The server failed to get SMI-S provider entity from database.](#)
- [HWN021802E Experienced SQL problems while working with database: The SQL error.](#)
- [HWN021803W The server did not get userid and or password for SMI-S provider The Service URL of the SMI-S provider from database.](#)
- [HWN021804E The server failed to access slp attributes for SMI-S provider The Service URL of the SMI-S provider from database.](#)
- [HWN021805E CIMOMManager failed to get a database mapper of type The type of the database mapper.](#)
- [HWN021806E CIMOMManager failed to get a valid mapper result from The type of the database mapper.](#)
- [HWN021807E CIMOMManager failed to get a proxy for calling slp discovery.](#)
- [HWN021808E The device cannot be contacted through any of the following SMI-S providers The comma separated list of IP and port for the SMI-S providers.. Possible causes are that the SMI-S providers are not accessible or the device is disconnected from the SMI-S providers.](#)
- [HWN021809E The host for SMI-S provider The service URL of the SMI-S providers. was not resolvable in DNS.](#)
- [HWN021810E The service URL for SMI-S provider The service URL of the SMI-S providers. is not valid.](#)
- [HWN021811I The operational status for device The ID of the device. on SMI-S provider The service URL of the SMI-S provider. has this value The operational status vector..](#)
- [HWN021812E The operational status for device The ID of the device. on SMI-S provider The service URL of the SMI-S provider. could not be retrieved because SMI-S provider is in status The SMI-S provider connection status..](#)
- [HWN021813E Fabric ID The ID of the fabric was not found.](#)
- [HWN021814E The device device id cannot be contacted through the SMI-S provider SMI-S provider service URL.](#)
- [HWN021899E Switch The wwn of the switch. has no associated Fabric.](#)
- [HWN021901E The virtual disk size cannot exceed maximum size when creating space efficient virtual disks.](#)
- [HWN021902E Invalid grain size. Valid values are valid values.](#)
- [HWN021903E Authentication to ip or name of host failed. Please specify correct authentication information.](#)
- [HWN021904E Connection to IP address or name of host failed with following operating system exception: exception text . Please make sure IP address is correct and machine is up and running. If this is a SVC V4 machine, it could be that its RAS interface is not up. If this is a SVC V5, make sure the SMI-S provider is up and running.](#)
- [HWN021905E Connection to IP address or name of host failed with following operating system exception: exception text .](#)
- [HWN021906E Failed to get native API entity from database.](#)
- [HWN021907E The IP address The service URL of the SMI-S providers. was not resolvable in DNS.](#)
- [HWN021908E Failed to get a proxy for calling NAPI discovery.](#)
- [HWN021909E There are no IO Groups available for Virtual Disk creation.](#)
- [HWN021910E Managed Disk ID The ID of the MDisk is not in unmanaged mode and cannot be added to the specified managed-disk group.](#)
- [HWN021911E Another probe of storage subsystem The Name+Nameformat of the storage subsystem is already in progress.](#)
- [HWN021912E Other probes of storage subsystems The list of Name+Nameformat of the storage subsystems are already in progress.](#)
- [HWN021913E IBM Spectrum Control Device Server could not write to directory The directory.](#)
- [HWN021914E SSH key file The SSH key file name is still in use, so it cannot be deleted.](#)
- [HWN021915E IBM Spectrum Control Device Server could not delete the file The file.](#)
- [HWN021916E The storage subsystem subsystem ID is not configured for file level management.](#)
- [HWN021917E An invalid parameter Name of the parameter which was invalid was specified. The corresponding file system mount point does not exist.](#)
- [HWN021919E The cluster ID The ID of the cluster. was not found.](#)
- [HWN021920E The export ID The ID of the export. was not found.](#)

- [HWN021921E The specified activity or protocol could not be used to change the export The ID of the export..](#)
- [HWN021922E The file system ID file_system_ID was not found.](#)
- [HWN021923E Invalid parameter Name of the parameter which was invalid. File system does not exist.](#)
- [HWN021924E The parameter Name of the parameter which was invalid is not a valid parameter.](#)
- [HWN021925E The fileset ID fileset_ID was not found.](#)
- [HWN021926E The WAN-cache source ID WAN_cache_source_id was not found.](#)
- [HWN021927E The WAN-cache ID WAN_cache_source_id was not found.](#)
- [HWN023000I The Optimization Execution task has started.](#)
- [HWN023001E The task to optimize the volumes was not completed successfully.](#)
- [HWN023002I The Optimization Execution task has completed.](#)
- [HWN023003I The Optimization Execution task retrieved number recommendations](#)
- [HWN023004I The Optimization Automation request persisted recommendations to be processed.](#)
- [HWN023005I The Optimization Execution task updated the status of number recommendations.](#)
- [HWN023006I The Optimization Automation request begins processing number recommendations.](#)
- [HWN023007W The recommendation being processed contains a virtual disk that is no longer detected.](#)
- [HWN023008W The recommendation for virtual disk vdisk name contains a source storage pool that is no longer detected.](#)
- [HWN023009W The recommendation for virtual disk vdisk name contains a target storage pool that is no longer detected.](#)
- [HWN023010I Virtual disk vdisk name was successfully migrated from storage pool source pool name to storage pool target pool name.](#)
- [HWN023011W The recommendation for virtual disk vdisk name contains a virtual disk that does not exist in the source storage pool source pool name or the target storage pool target pool name.](#)
- [HWN023012W The recommendation for virtual disk vdisk name contains a non-mirrored virtual disk that is now a mirrored virtual disk.](#)
- [HWN023013W The recommendation for virtual disk vdisk name contains a mirrored virtual disk that is now a non-mirrored virtual disk.](#)
- [HWN023014I The recommendation for virtual disk vdisk name requires more space on target pool target pool name to be processed.](#)
- [HWN023015I Virtual disk vdisk name will now be migrated from storage pool source pool name to storage pool target pool name.](#)
- [HWN023016I Successfully added virtual disk copy to virtual disk vdisk name.](#)
- [HWN023017I Synchronization for virtual disk vdisk name has completed synchronization percent% and requires about seconds to complete seconds to complete.](#)
- [HWN023018I Synchronization for virtual disk vdisk name has completed.](#)
- [HWN023019I Successfully removed a virtual disk copy from virtual disk vdisk name.](#)
- [HWN023020I Successfully changed the synchronization rate of virtual disk vdisk name to syncrate%.](#)
- [HWN023021I Successfully changed the primary copy of virtual disk vdisk name.](#)
- [HWN023022E There is no space available on target pool target pool name to migrate the virtual disk vdisk name.](#)
- [HWN023023E Unable to submit request to add vdisk copy command for virtual disk vdisk name due to rc \(rc\).](#)
- [HWN023024E Unable to complete request to add vdisk copy command for virtual disk vdisk name due to rc \(rc\).](#)
- [HWN023025E Unable to submit request to get vdisk synchronization progress for virtual disk vdisk name due to rc \(rc\).](#)
- [HWN023026E Unable to complete request to get vdisk synchronization progress for virtual disk vdisk name due to rc \(rc\).](#)
- [HWN023027E Unable to submit request to remove vdisk copy command for virtual disk vdisk name due to rc \(rc\).](#)
- [HWN023028E Unable to complete request to remove vdisk copy command for virtual disk vdisk name due to rc \(rc\).](#)
- [HWN023029E Unable to submit request to change the synchronization rate for virtual disk vdisk name due to rc \(rc\).](#)
- [HWN023030E Unable to complete request to change the synchronization rate for virtual disk vdisk name due to rc \(rc\).](#)
- [HWN023031E Unable to submit request to change the primary copy for virtual disk vdisk name due to rc \(rc\).](#)
- [HWN023032E Unable to complete request to change the primary copy for virtual disk vdisk name due to rc \(rc\).](#)
- [HWN023033E The request failed. Message from failed request: message.](#)
- [HWN023034E The Optimization Automation job completed with errors in the recommendations.](#)
- [HWN023035W The Optimization Execution task completed with warnings.](#)
- [HWN023036E The request failed because there were not enough extents in the storage pool.](#)
- [HWN023037E The request failed because the number of copies of this volume would exceed the limit.](#)
- [HWN023038E The request failed because the copy specified does not exist.](#)
- [HWN023039E The following exception occurred during a migration request: exception](#)
- [HWN023040E The migration request for volume vdisk name is already being processed.](#)
- [HWN023041W The request to migrate the mirrored volume vdisk name is suspended because the secondary volume is offline.](#)
- [HWN023042E The secondary copy needed for migration does not exist.](#)
- [HWN023043I The mirrored volume migration for volume vdisk name will be ignored.](#)
- [HWN023044I The mirrored volume migration for volume vdisk name will result in the current secondary volume becoming the primary volume.](#)
- [HWN023045I The mirrored volume migration for volume vdisk name will result in the primary volume being migrated to the target pool.](#)
- [HWN023046I The Migration of the previously abandoned Optimization Automation job has started.](#)
- [HWN023047I The Migration of the previously abandoned Optimization Automation job has completed.](#)
- [HWN023048I The Optimization Automation cancellation job jobname has started.](#)
- [HWN023049E The Optimization Automation cancellation job completed with errors.](#)
- [HWN023050I The Optimization Automation cancellation job jobname has completed.](#)
- [HWN023051I The Optimization Automation job jobname will be canceled.](#)
- [HWN023052W The Optimization Automation job is not in progress.](#)
- [HWN023053I The migration of volume vdisk name has been canceled.](#)
- [HWN023054W The Optimization Automation job was canceled.](#)
- [HWN023055I The volume that was chosen for transformation, vdisk name, is a secondary volume in a mirrored volume relationship. The secondary volume will be migrated to the specified target pool or converted as specified.](#)
- [HWN024000I An optimization analysis task was started.](#)
- [HWN024001I The analysis is completed.](#)
- [HWN024002W Unable to retrieve any policy for Tier value.](#)
- [HWN024003I Analyzed number of volumes volumes on tier tier_number for storage virtualizer subsystem_name.](#)
- [HWN024006W No target pools in subsystem value were selected.](#)
- [HWN024011W Destination storage pool value in subsystem value was not considered. Reason: value.](#)
- [HWN024012I It is recommended that number of volumes volumes on tier source_tier_number are moved to tier target_tier_number.](#)
- [HWN024015I The optimization analysis of the value subsystem was started.](#)
- [HWN024016W Volume value is already in the destination storage pool value. No recommendations will be generated for the volume.](#)
- [HWN024018W No destination storage pools in Tier value have been specified for subsystem value.](#)
- [HWN024019W The following pools on tier tier_number on the storage_system storage system cannot be balanced by redistributing or re-tiering volumes: pool_names.](#)
- [HWN024020I Started analysis to balance pools on tier value.](#)
- [HWN024021W The pool_name pool on tier tier_number on the storage_system storage system cannot be balanced by redistributing the volumes.](#)

- [HWN024027I Storage Pool pool name has insufficient available space for volume volume name in storage pool pool name.](#)
- [HWN024030W One or more entities specified as input for the analysis could not be found or pools or volumes in some input entities could not be found.](#)
- [HWN024031W One or more entities specified as candidate destinations for the analysis could not be found.](#)
- [HWN024032W For one or more mirrored volumes, both the primary and the secondary volume copies were chosen for transformation. You cannot transform both volume copies in the same transform task. Only the primary volume copies are included for transformation. You can transform the secondary volume copies in a separate transformation.](#)
- [HWN024033W The volume volume name cannot be analyzed because it is not in a capacity pool.](#)
- [HWN024034W The pool pool name cannot be analyzed because the pool is not in a capacity pool.](#)
- [HWN024035W The storage virtualizer system name cannot be analyzed because the storage virtualizer is not in a capacity pool.](#)
- [HWN024036W The operation to transform the volumes on the subsystem name storage virtualizer cannot be completed because the destination pools were not available.](#)
- [HWN024037E An unexpected error occurred. The operation to transform the volumes on the subsystem name storage virtualizer cannot be completed because the destination pools were not identified.](#)
- [HWN024043I The capacity pools of the source volumes were selected as the target pools.](#)
- [HWN024046I The option that was selected to handle volumes with mirrored volumes is: After optimization, set the copy of the secondary volume in the destination pool as the primary volume. The original secondary volume remains the secondary volume.](#)
- [HWN024047I The number of days for collecting performance data to analyze the volumes is set to performance_data_collection_period.](#)
- [HWN024050I Automatic tiering was selected to tier the volumes.](#)
- [HWN024051I The tiering analysis is starting.](#)
- [HWN024052I Tier tier# has an I/O density threshold value of value per second per GiB.](#)
- [HWN024053I Tier tier#, has a file age threshold value of value percent of files last accessed within time unit.](#)
- [HWN024054I The real capacity for the thin-provisioned volumes is set to value unit.](#)
- [HWN024055I The auto expand property of the thin-provisioned volumes is set to yes/no.](#)
- [HWN024056I The warning level for thin-provisioned volumes is set to value %.](#)
- [HWN024057I The grain size that was specified for the thin-provisioned volumes is grain_size KiB.](#)
- [HWN024058I The real capacity for the compressed volumes is set to value unit.](#)
- [HWN024059I The auto expand property for the compressed volumes is set to yes/no.](#)
- [HWN024060I The warning level for the compressed volumes is set to value.](#)
- [HWN024061I The option that was selected to handle volumes with mirrored volumes is: After optimization, set the secondary volume as the primary volume. The volume in the destination pool is the secondary volume.](#)
- [HWN024062I The option that was selected for mirrored volumes is: Do not optimize volumes with mirrored volumes.](#)
- [HWN024066I Tier tier# has an I/O rate threshold value of value I/O per second.](#)
- [HWN024067W Recommendations cannot be generated for number_of_volumes volumes because the volumes do not meet the tiering criteria for tier current_tier_number or for any lower tier.](#)
- [HWN024068W Recommendations cannot be generated to move number_of_volumes volumes from source_tier to tier target_tier_number due to the pool activity limit value.](#)
- [HWN024069W Recommendations cannot be generated to move number_of_volumes volumes from tier source_tier to tier target_tier_number because the destination storage_pools do not have enough space.](#)
- [HWN024070I The analysis to optimize subsystem storage_subsystem was completed.](#)
- [HWN024071I The option that was selected was to restrict the placement of volumes in capacity_pools to destination storage_pools in the same capacity pool.](#)
- [HWN024072W No file age information for volume volume name.](#)
- [HWN024073W Storage pool {0} in tier {1} needs at least one additional storage pool in the same tier for the Balance Analysis to run on this tier.](#)
- [HWN024074W Storage pool {0} in tier {1} and capacity pool {2} needs at least one additional storage pool in the same tier and capacity pool for the Balance Analysis to run within this capacity pool and on this tier.](#)
- [HWN024075W number_of_volumes volumes from storage pool pool could not be moved to the destination storage_pools because the destination storage_pools do not have enough space.](#)
- [HWN024076W number_of_volumes volumes from storage pool pool could not be moved to the destination_pools because the destination storage_pools are not in the same capacity pool.](#)
- [HWN024077W number_of_volumes volumes from storage pool pool could not be moved to the destination storage_pools because the destination storage_pools would have exceeded the pool activity limit value.](#)
- [HWN024078W number_of_volumes volumes from storage pool pool could not be moved to the destination storage_pools because the destination storage_pools already have a volume copy.](#)
- [HWN024079W Because of an internal error, the number of volumes in the pool storage pool that could not be moved to destination storage_pools is number_of_volumes.](#)
- [HWN024080W Destination storage pool pool already contains a copy of storage volume volume.](#)
- [HWN024081W Because the destination storage pool does not have sufficient available space, the volume storage volume in the source_pool storage pool cannot be moved to the destination_pool destination storage pool.](#)
- [HWN024082W Because the destination storage pool contains a copy of the mirrored volume, the volume volume in the source_pool storage pool cannot be moved to the destination_pool destination storage pool.](#)
- [HWN024083W Because of an internal error, the volume storage volume in the spool storage pool could not be moved to the destination_pool destination storage pool.](#)
- [HWN024084W Because the destination storage_pools contain one or more copies of the mirrored volumes, the number of volumes that could not be moved from tier source_tier to tier target_tier is number_of_volumes.](#)
- [HWN024085W The pool_name storage pool cannot be balanced because the tier level of the pool was reset to none.](#)
- [HWN024086E Recommendations cannot be generated because the tier level of the destination_pool_name destination storage pool was reset to none.](#)
- [HWN024087W Recommendations cannot be generated for one or more of the volumes because collocated volumes cannot be placed in the same destination storage pool.](#)
- [HWN024088I The option to collocate volumes that are assigned to the same server or hypervisor was selected.](#)
- [HWN024089I The option to collocate volumes that are assigned to the same server or hypervisor was not selected.](#)
- [HWN024090W Because the storage_pools do not meet the service class requirements, the number of volumes that cannot be moved is no_volumes.](#)
- [HWN024091W If the recommendation to move the volume_name volume to the storage_pool_name storage pool is implemented, the service class requirements of the volume_name volume cannot be met.](#)
- [HWN024092W Recommendations cannot be generated to move number_of_volumes volumes from tier source_tier to tier target_tier_number because the destination storage_pools do not meet the service class requirements of the volumes.](#)
- [HWN024093I The number of volumes on tier tier_level that were not analyzed because of the instruction to exclude mirrored volumes from the analysis is number_of_volumes volumes.](#)
- [HWN024094W Valid target pools were not selected for the subsystem name storage virtualizer.](#)
- [HWN024095I The grain size for the thin-provisioned volumes was set to the default value of grain_size KiB.](#)

- [HWN024096W Volumes in the pool_name pool on tier tier_level cannot be moved to a higher tier to reduce the activity level of the pool to the user-defined level.](#)
- [HWN024097W Volumes in the pool_name pool on tier tier_level cannot be moved to a lower tier to reduce the activity level of the pool to the user-defined level.](#)
- [HWN024098W Cannot generate recommendations to tier volumes from the storage_system_name storage system because all of the source volumes are in the selected destination storage pools.](#)
- [HWN024099I The number of volumes that were excluded from the analysis to plan the tiering of the storage_system_name storage system is vols_count. The volumes were excluded because performance data is not available for the volumes.](#)
- [HWN024100I The number of volumes that were excluded from the analysis to plan the tiering of the storage_system_name storage system is vols_count. The volumes were excluded from the analysis because the capacity of the volumes is zero.](#)
- [HWN024101I The number of volumes that were excluded from the analysis to plan the tiering of the storage_system_name storage system is vols_count. The volumes were excluded from the analysis because the volumes are not assigned to pools that are tiered or the thresholds were not defined for the tiers.](#)
- [HWN024102W The recommendation to move the storage_volume_name volume from the source_pool_name storage pool to the target_pool_name storage pool was not generated because the status of the destination pool is offline or excluded.](#)
- [HWN024103I Reclaiming volumes](#)
- [HWN024104I Planning for tiering volumes](#)
- [HWN024105W The recommendation to move the storage_volume_name volume from the source_pool_name storage pool to the target_pool_name storage pool will not be executed because the status of the destination pool is offline or excluded.](#)
- [HWN024106W The recommendation to move the storage_volume_name volume from the source_pool_name storage pool was not generated because the status of the volume is offline.](#)
- [HWN024107W The recommendation to move the storage_volume_name volume from the source_pool_name storage pool to the target_pool_name storage pool will not be executed because the status of the volume is offline.](#)
- [HWN024108E The recommendations can't be shown because the analysis was not completed.](#)
- [HWN024109W The data for the previous analysis of the storage_subsystem storage system was not deleted.](#)
- [HWN024110E Volumes reclamation analysis failed for storage_subsystem storage subsystem.](#)
- [HWN024111W Recommendations cannot be generated to move number_of_volumes volumes from tier source_tier to tier target_tier number because there is no potential destination pool assigned to the recommended tier.](#)
- [HWN024112W Cannot generate recommendations to tier volumes from the storage_system_name storage system because the source storage pools and the selected destination storage pools are assigned to the same tier.](#)
- [HWN024200I The days of the week to include in the analysis: days_of_week.](#)
- [HWN024201I The time window for the performance data to include in the analysis is set to start time - end time.](#)
- [HWN024202I The time window for the performance data to include in the analysis is set to start time - end time. The end time occurs on the next day.](#)
- [HWN024203W The volume storage_volume_name cannot be converted or moved because the target pools do not have sufficient available space or the target pool types are incorrect for the operation.](#)
- [HWN025000I Storage pool value in storage system value has storage from different types of back-end storage systems. Back-end disk data cannot be determined.](#)
- [HWN025001I Storage pool value in storage system value has storage from unknown back-end storage system\(s\). Back-end disk data cannot be determined.](#)
- [HWN025002I Storage pool value in storage system value has storage from multiple back-end storage systems or from multiple pools in a single storage system. Back-end disk data cannot be determined.](#)
- [HWN025003I Storage pool value in storage system value has storage from a back-end storage pool with multiple disk types. Back-end disk data cannot be determined.](#)
- [HWN025004I Storage pool value in storage system value has storage from a back-end storage pool with a mixed raid type. Back-end disk data cannot be determined.](#)
- [HWN025005I Storage pool value in storage system value has storage from a back-end storage pool with multiple raid types. Back-end disk data cannot be determined.](#)
- [HWN025006I Storage pool value in storage system value has storage from back-end disks of unknown type. Back-end disk data cannot be determined.](#)
- [HWN025007I Storage pool value in storage system value has storage from unknown number of back-end disks. Back-end disk data cannot be determined.](#)
- [HWN025008I Storage pool value in storage system value has storage from back-end disks with unknown raid type. Back-end disk data cannot be determined.](#)
- [HWN025009E Connection to Data Server failed. Make sure Data Server is up.](#)
- [HWN025011W All of the target ports for the storage system are used for the provisioning request. The request might take a long amount of time.](#)
- [HWN025010I Collecting parent pool volumes for storage system: storage_system identification.](#)
- [HWN025011E The port the target port ID has a usage restriction which prevents it from being used as a target port for volume assignment.](#)
- [HWN025012E The invocation of CIM method ExposePaths failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method.](#)
- [HWN025013E The invocation of CIM method HidePaths failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method.](#)
- [HWN025014E The invocation of CIM method CreateOrModifyElementFromStoragePool failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method.](#)
- [HWN025015E The invocation of CIM method ReturnToStoragePool failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method.](#)
- [HWN025016E The invocation of CIM method DeleteStorageHardwareID failed on SMI-S provider Name of SMI-S provider . The return code is Return code of method.](#)
- [HWN025017E A CLI command failed. Check the logs from EP working dir.](#)
- [HWN025018E An error occurred when attempting to parse the file File name.](#)
- [HWN025019E The requested operation failed. Check the logs from EP working dir.](#)
- [HWN025020E The volume cannot be created. The volume of size Size in pool Pool on storage system Subsystem cannot be created. The pool might already have the maximum number of volumes allowed.](#)
- [HWN025021E Unable to resolve the address for the device because the request was not processed by the data collector.](#)
- [HWN025022E The data collection detected storage system New Subsystem with serial number new serial number instead of expected serial number expected serial number.](#)
- [HWN025025I Starting the task to send the report for schedule Schedule Id by email.](#)
- [HWN025026I The report title report is being created.](#)
- [HWN025027I The report title report with ID report id is being sent by email to the reports recipients.](#)
- [HWN025028I The report title report with ID report id was sent by email to the reports recipients.](#)
- [HWN025029E Can't retrieve the configured settings of the report for schedule Schedule Id .](#)
- [HWN025030E The report can't be sent because the email server was not configured.](#)
- [HWN025031E Can't send the report title report with ID report id by email because of the following error: reported_error.](#)
- [HWN025031I To view the report, choose HTML as the message format or use an email application that supports HTML message formats.](#)
- [HWN025032E Job failed during post processing of collected data from the data source.](#)
- [HWN025033E Failed to send the report name report for schedule Schedule Id.](#)
- [HWN025034I Created number_of_servers agentless servers automatically.](#)

- [HWN025035I Removed number of servers agentless servers automatically.](#)
- [HWN025036E Can't save the report in the directory.](#)
- [HWN025037E Can't save the report because the path specifies a file name instead of a directory name.](#)
- [HWN025038E Can't save the report, because the directory doesn't exist.](#)
- [HWN025039E Can't save the report because the directory doesn't have enough disk space.](#)
- [HWN025040I The report title report with ID report id is being saved as report file name in the full path directory.](#)
- [HWN025041I The report title report with ID report id was saved as report file name in the full path directory.](#)
- [HWN099990I The method name of the Device Server method of the device server returned return value @\(execution context information\).](#)
- [HWN099991I info trace message@\(execution context information\)](#)
- [HWN099992W warning trace message@\(execution context information\)](#)
- [HWN099993E error/exception trace message @\(execution context information\)](#)
- [HWN099994I An object of class name of the class has been instantiated @\(execution context information\).](#)
- [HWN099995I |=== class name.method name entry, parameter\(s\); parameter value\(s\) @\(execution context information\).](#)
- [HWN099996I ===| class name.method name exit, return value: method return value \(execution time in milliseconds\) @\(execution context information\).](#)
- [HWN099997I External service name of the \(DM\) external service will be invoked with parameter\(s\) parameter value\(s\)@\(execution context information\).](#)
- [HWN099998I Invocation of external service name of the \(DM\) external service returned result invocation result@\(execution context information\).](#)
- [HWN099999I The method name of the device server method of the device server was invoked with parameters invocation parameters@\(execution context information\).](#)
- [HWN200000I Probe of switch switch_name completed successfully.](#)
- [HWN200001I Started post-processing tasks after data was collected for switch switch_name.](#)
- [HWN6001I Operation operation completed successfully.](#)
- [HWN6002I Unable to set up NLS message file processing.](#)
- [HWN6003E Unable to set up tracing.](#)
- [HWN6004E Operation operation failed.](#)
- [HWN6005E Unknown operation operation.](#)
- [HWN6006E Could not initialize connection, rc is rc](#)
- [HWN6007E Could not parse command arguments: arg](#)
- [HWN6008E Error processing command: command](#)
- [HWN6009E Missing 'operation' property in input file](#)
- [HWN6010I Task arg completed successfully](#)
- [HWN6011E Task arg failed](#)
- [HWN6012E Cannot connect to this IP, switching to IP](#)
- [HWN6013E An IBM XIV CLI command failed. The error is arg.](#)
- [HWN6014I Command arg completed successfully.](#)
- [HWN6015E Command command failed.](#)
- [HWN6016I Connected with IP address IP](#)
- [HWN6017I Started creation of volume with size size in pool pool.](#)
- [HWN6018I Volume creation completed successfully. New volume volume created with size size in pool pool.](#)
- [HWN6019I Started deletion of volume volume in pool pool.](#)
- [HWN6020I Volume deletion completed successfully. Volume volume deleted in pool pool](#)
- [HWN6021I Started creation of host host with initiator ports ports](#)
- [HWN6022I Finished creation of host host with initiator ports ports](#)
- [HWN6023I Started assignment of volume volume to host host.](#)
- [HWN6024I Finished assignment of volume volume to host host.](#)
- [HWN6025I Started unassignment of volume volume from host host.](#)
- [HWN6026I Finished unassignment of volume volume from host host](#)
- [HWNEP0001I Successfully persisted number of count instances.](#)
- [HWNEP0002E The probe failed as the data collector couldn't write to its output file, value.](#)
- [HWNEP0003E A DS8000 ESSNI command failed. The error code is error_code.](#)
- [HWNEP0004I Started creation of volume group volume_group.](#)
- [HWNEP0005I Finished creation of volume group volume_group with subsystem volume_group number number.](#)
- [HWNEP0006I Started adding volumes, with serial numbers volume_list, to subsystem volume_group volume_group_number.](#)
- [HWNEP0007I Finished adding volumes to volume_group.](#)
- [HWNEP0008I Started assignment of host host on subsystem subsystem to volume_group volume_group.](#)
- [HWNEP0009I Finished assigning host on subsystem subsystem to volume_group volume_group.](#)
- [HWNEP0010I Started removing volumes, with serial numbers volume_list, from subsystem volume_group volume_group_number.](#)
- [HWNEP0011I Finished removing volumes, with serial numbers volume_list, from subsystem volume_group volume_group_number.](#)
- [HWNEP0012I Increased virtual capacity of storage pool storage_pool on subsystem subsystem to size size.](#)
- [HWNEP0013I Collecting pools for storage system storage_system identification.](#)
- [HWNEP0014I Collecting volumes for lss logical subsystems on storage system storage_system identification.](#)
- [HWNEP0015I Collecting volume groups on storage system storage_system identification.](#)
- [HWNEP0016I Collecting hosts on storage system storage_system identification.](#)
- [HWNEP0017I value Hosts Found.](#)
- [HWNEP0018I Launching external process for devices devices.](#)
- [HWNEP0019I External process for devices devices completed successfully.](#)
- [HWNEP0020E Could not create connection to NAPI The IP for the NAPI..](#)
- [HWNEP0021E ESSNI API query for Space Efficient Volume failed with ESSNI code ESSNI Code. Data from ESSNI is considered suspect.](#)
- [HWNEP0022I Started deletion of volume_group with number volume_group_number.](#)
- [HWNEP0023I Finished deletion of volume_group with number volume_group_number.](#)
- [HWNEP0100I Probing Volumes for Storage System: value](#)
- [HWNEP0101I Number of Volumes currently found: value. Continuing to probe Volumes.](#)
- [HWNEP0102I value Volumes found.](#)
- [HWNEP0103I Probing Configured Disks for Storage System: value.](#)
- [HWNEP0104I Number of Configured Disks Found Currently: value. Continuing to Probe Disks.](#)
- [HWNEP0105I value Configured Disks Found.](#)
- [HWNEP0106I Probing Views of Host Initiator access to Volumes.](#)
- [HWNEP0107I Finished probing Views.](#)
- [HWNEP0108I Initializing Probe for storage system storage_system identification.](#)
- [HWNEP0109I Probe for storage system storage_system identification completed.](#)
- [HWNEP0110I Collecting Nodes and fibre channel ports for storage system storage_system identification.](#)

- [HWNEP0111E The connection to the storage device failed. The error code is error_code.](#)
- [HWNEP0113E The cluster IP address is not specified in the configuration file.](#)
- [HWNEP0114E The trustore location is not specified in the configuration file.](#)
- [HWNEP0115E The IBM Spectrum Control data is out of synch with the device configuration and a re-probe is required for device device_name.](#)
- [HWNEP0116E The user configured for the subsystem subsystem_name is not permitted to perform the requested action.](#)
- [HWNEP0117E The virtual disk \(VDisk\)-to-host mapping was not created because the volume vdiskName is already mapped to the hostName host for the Device deviceName](#)
- [HWNEP0115I Starting Control Process for storage system storage_system identification.](#)
- [HWNEP0116I Started deletion of volume VolumeID on subsystem Subsystem.](#)
- [HWNEP0117I Volume deletion completed successfully. Volume VolumeID on subsystem Subsystem was deleted.](#)
- [HWNEP0118I Started adding Managed Disk\(s\) Managed Disk ID to Managed-disk group Managed Disk group name on subsystem Subsystem.](#)
- [HWNEP0119I Finished adding Managed Disk\(s\) Managed Disk ID to Managed-disk group Managed Disk group name on subsystem Subsystem.](#)
- [HWNEP0120I Started creation of volume with size Size in pool Pool on subsystem Subsystem](#)
- [HWNEP0121I Volume creation completed successfully. New volume VolumeID created with size Size in pool Pool on subsystem Subsystem.](#)
- [HWNEP0122I Started assignment of volume VolumeID on subsystem Subsystem to initiator port Initiator Port on host Host.](#)
- [HWNEP0123I Finished assignment of volume VolumeID on subsystem Subsystem to initiator port Initiator Port on host Host Name.](#)
- [HWNEP0124I Started unassignment of volume VolumeID on subsystem Subsystem from initiator port Initiator Port on host Host Name.](#)
- [HWNEP0125I Finished unassignment of volume VolumeID on subsystem Subsystem from initiator port Initiator Port on host Host Name.](#)
- [HWNEP0126I Started creation of host host name on subsystem Subsystem with initiator ports WWPNs.](#)
- [HWNEP0127I Finished creation of host host name on subsystem Subsystem with initiator ports WWPNs.](#)
- [HWNEP0128I Host name hostName already exists for the WWPNs wwpgns on the device Subsystem](#)
- [HWNEP0129E The operation failed because the device returned unexpected values.](#)
- [HWNEP0130E A IBM XIV CLI command failed. The error is error_code.](#)
- [HWNEP0131I The host definition for host host name on subsystem Subsystem contains additional Hostports WWPNs that will also be assigned to Volume VolumeID.](#)
- [HWNEP0132E The unassignment of Volume VolumeID from hostport WWPN failed because the definition for host host name on subsystem Subsystem contains additional hostports WWPNs.](#)
- [HWNEP0133E Error invoking the external process for device device_name.](#)
- [HWNEP0134E Following exception occurred: exception.](#)
- [HWNEP0135E External process failed with error code error_code.](#)
- [HWNEP0136E Error connecting to IP address with user ID user ID.](#)
- [HWNEP0137I Job job ID submitted for device device_name.](#)
- [HWNEP0138I External process was successfully executed for device device_name.](#)
- [HWNEP0139I An instruction was issued to add a copy of the volume_name volume_size-byte volume in the pool_name pool on the storage_system_name storage system.](#)
- [HWNEP0140I The copy of the volume_name volume_size-byte volume with the copy ID of VolumeID in the pool_name pool on the storage_system_name storage system was added successfully.](#)
- [HWNEP0141I Probing Internal Drives for Storage System: value.](#)
- [HWNEP0142I Number of Internal Drives Found Currently: value. Continuing to Probe Internal Drives.](#)
- [HWNEP0143I value Internal Drives Found.](#)
- [HWNEP0144I Probing Pools for Storage System: value.](#)
- [HWNEP0145I Number of Pools Found Currently: value. Continuing to Probe Pools.](#)
- [HWNEP0146I value Pools Found.](#)
- [HWNEP0147I Collecting asset and status information about storage_system_id storage system.](#)
- [HWNEP0148I Collecting cluster information for storage_system_id storage system.](#)
- [HWNEP0149I Collecting file system exports for storage_system_id storage system.](#)
- [HWNEP0150I Collecting nodes for storage_system_id storage system.](#)
- [HWNEP0151I Collecting file systems for storage_system_id storage system.](#)
- [HWNEP0152I Collecting pools for storage_system_id storage system.](#)
- [HWNEP0153I Collecting file system storage for storage_system_id storage system.](#)
- [HWNEP0154I Collecting filesets for storage_system_id storage system.](#)
- [HWNEP0155I Collecting links between file systems and nodes for storage_system_id storage system.](#)
- [HWNEP0156I Collecting quotas for storage_system_id storage system.](#)
- [HWNEP0157I Collecting file system snapshots for storage_system_id storage system.](#)
- [HWNEP0158I Collecting capacity for file_system_id file system.](#)
- [HWNEP0159I Creating the export export name on cluster cluster_name.](#)
- [HWNEP0160I The export export name on cluster cluster_name with path export path was created.](#)
- [HWNEP0161I The export export name on cluster cluster_name is being changed.](#)
- [HWNEP0162I The export export name on cluster cluster_name was changed.](#)
- [HWNEP0163I Setting quota quota type - quota name on file system file system_name.](#)
- [HWNEP0164I Quota quota type - quota name on file system file system_name has been created.](#)
- [HWNEP0165I Checking quota on file system file system_name.](#)
- [HWNEP0166I Quota on file system file system_name has been checked.](#)
- [HWNEP0167I The export export name on cluster cluster_name is being removed.](#)
- [HWNEP0168I The export export name on cluster cluster_name was removed.](#)
- [HWNEP0169E Command: command did not complete. IBM SONAS CLI message](#)
- [HWNEP0170I Creating fileset fileset name on file system files system_name.](#)
- [HWNEP0171I Successfully created fileset fileset name on file system file system_name.](#)
- [HWNEP0172I Removing fileset fileset name on file system files system_name.](#)
- [HWNEP0173I Successfully removed fileset fileset name on file system file system_name.](#)
- [HWNEP0174I Modifying fileset fileset name on file system files system_name.](#)
- [HWNEP0175I Successfully modified fileset fileset name on file system file system_name.](#)
- [HWNEP0176I Creating file system file system on cluster cluster_name.](#)
- [HWNEP0177I Successfully created file system file system on cluster cluster_name.](#)
- [HWNEP0178I Changing file system file system on cluster cluster_name.](#)
- [HWNEP0179I Successfully changed file system file system on cluster cluster_name.](#)
- [HWNEP0180I Removing file system file system on cluster cluster_name.](#)
- [HWNEP0181I Successfully removed file system file system on cluster cluster_name.](#)
- [HWNEP0182I Mounting file system file system.](#)
- [HWNEP0183I Successfully mounted file system file system.](#)

- [HWNEP0184I Unmounting file system file system .](#)
- [HWNEP0185I Successfully unmounted file system file system .](#)
- [HWNEP0186I Linking fileset fileset on file system file system .](#)
- [HWNEP0187I Successfully linked fileset fileset on file system file system .](#)
- [HWNEP0188I Unlinking fileset fileset on file system file system .](#)
- [HWNEP0189I Successfully unlinked fileset fileset on file system file system .](#)
- [HWNEP0190E The IBM Spectrum Control server could not connect to IP address using the SSH protocol.](#)
- [HWNEP0191E The IBM Spectrum Control server could not authenticate with IP address using the SSH protocol.](#)
- [HWNEP0192E The IBM Spectrum Control server could not execute a command on the IBM Storwize V7000 Unified/IBM SONAS device at IP address .](#)
- [HWNEP0193E The command name command failed because the following command executed on the NAS device failed with the return code return code : command returned: command output](#)
- [HWNEP0195I modify fileset](#)
- [HWNEP0196I change export](#)
- [HWNEP0197I create export](#)
- [HWNEP0198I remove export](#)
- [HWNEP0199I create fileset](#)
- [HWNEP0200I link fileset](#)
- [HWNEP0201I remove fileset](#)
- [HWNEP0202I unlink fileset](#)
- [HWNEP0203I change filesystem](#)
- [HWNEP0204I create filesystem](#)
- [HWNEP0205I mount filesystem](#)
- [HWNEP0206I remove filesystem](#)
- [HWNEP0207I unmount filesystem](#)
- [HWNEP0208I check quota](#)
- [HWNEP0209I set quota](#)
- [HWNEP0210I probe](#)
- [HWNEP0211W The command name command completed , however during post-processing the following command executed on the NAS device failed with the return code return code : command returned: command output As a result, the IBM Spectrum Control database is now out of sync with the current state of the NAS device.](#)
- [HWNEP0212I create disk in modifying file system](#)
- [HWNEP0213I Started deletion of host host name on subsystem Subsystem .](#)
- [HWNEP0214I Finished deletion of host host name on subsystem Subsystem .](#)
- [HWNEP0215I Collecting cache information for storage_system_id storage system.](#)
- [HWNEP0216I remove cached source](#)
- [HWNEP0217I create cached node](#)
- [HWNEP0218I remove cached node](#)
- [HWNEP0219I create cache](#)
- [HWNEP0220I remove cache](#)
- [HWNEP0221I modify cache source](#)
- [HWNEP0222I Creating cache source cache_source_name on cluster file_system_name.](#)
- [HWNEP0223I Created cache source cache_source_name on cluster file_system_name.](#)
- [HWNEP0224I Removing cache source cache_source_name on cluster file_system_name.](#)
- [HWNEP0225I Removed cache source cache_source_name on cluster file_system_name.](#)
- [HWNEP0226I Modifying cache source cache_source_name on cluster file_system_name.](#)
- [HWNEP0227I Modified cache source cache_source_name on cluster file_system_name.](#)
- [HWNEP0228I Creating cache cache_name on file system file_system_name.](#)
- [HWNEP0229I Created cache cache_name on file system file_system_name.](#)
- [HWNEP0230I Removing cache cache_name on file system file_system_name.](#)
- [HWNEP0231I Removed cache cache_name on file system file_system_name.](#)
- [HWNEP0232I Modifying cache cache_name on file system file_system_name.](#)
- [HWNEP0233I Modified cache cache_name on file system file_system_name.](#)
- [HWNEP0234I modify cache](#)
- [HWNEP0235I create cached source](#)
- [HWNEP0236I Configuring nodes node_names as cached nodes.](#)
- [HWNEP0237I Configured nodes node_names as cached nodes.](#)
- [HWNEP0238I Unconfiguring cached nodes node_names.](#)
- [HWNEP0239I Unconfigured cached nodes node_names.](#)
- [HWNEP0240I Executed control operation on cache cache_name on filesystem file_system_name .](#)
- [HWNEP0241I control cache](#)
- [HWNEP0242I run prepop](#)
- [HWNEP0243I list prepop](#)
- [HWNEP0244I Retrieving cache prepopulation status for file system file_system_name .](#)
- [HWNEP0245I Cache prepopulation status for file system file_system_name has been retrieved.](#)
- [HWNEP0246I Prepopulate cache data for fileset fileset_name on file system file_system_name using policy policy_name.](#)
- [HWNEP0247I Command to pre populate cached data for fileset fileset_name was successful.](#)
- [HWNEP0248W An error was encountered while parsing protocol options for export export_name. The options were not persisted, the probe will continue.](#)
- [HWNEP0249W The connection to the storage device failed. The error code is error_code.](#)
- [HWNEP0250I Started adding initiator port\(s\) initiator ports to host host name on subsystem subsystem .](#)
- [HWNEP0251I Finished adding initiator port\(s\) initiator ports to host host name on subsystem subsystem .](#)
- [HWNEP0252W A CLI command completed with warning. The warning message is : warning_message](#)
- [HWNEP0253W Volume creation completed with warning. New volume VolumeID created with size Size in pool Pool on subsystem Subsystem .](#)
- [HWNEP0254W Volume deletion completed with warning. Volume VolumeID on subsystem Subsystem was deleted.](#)
- [HWNEP0255I The task to execute the recommendations for optimizing the volumes on the storage system with an ID of storage_system_id was paused.](#)
- [HWNEP0256I The task for optimizing the volumes on the storage system with an ID of storage_system_id was canceled.](#)
- [HWNEP0257I The task for optimizing the volumes on the storage system with an ID of storage_system_id was resumed.](#)
- [HWNEP0258E The optimization task cannot be paused because the synchronization rate for the volume cannot be reset. The ID of the volume is volume_id and the ID of the storage system is storage_system_id.](#)
- [HWNEP0259E The optimization task cannot be resumed because the synchronization rate for the volume cannot be reset. The ID of the volume is volume_id and the ID of the storage system is storage_system_id.](#)

- [HWNEP0260I Started creation of host port host port name on storage system Storage System with initiator port WWPN.](#)
- [HWNEP0261I Finished creation of host port host port name on storage system Storage System with initiator port WWPN.](#)
- [HWNEP0262E The recommendation for the volume_name volume was not implemented because the command that was issued by the storage virtualizer returned the following error: error_message](#)
- [HWNEP0263I The synchronization of the volume_name volume with the volume copy was successful.](#)
- [HWNEP0264E The synchronization of the volume_name volume with the volume copy was unsuccessful.](#)
- [HWNEP0265E The CLI command that was issued for the storage_system_name storage system failed and generated the following error: error_message](#)
- [HWNEP0266I Started expanding the capacity of volume volume on subsystem subsystem from oldsize to newsize bytes.](#)
- [HWNEP0267I Finished expanding the capacity of volume volume on subsystem subsystem to newsize bytes.](#)
- [HWNEP0268E The server operating system or version is not supported by IBM Spectrum Control for IBM Spectrum Scale.](#)
- [HWNEP0269E The IBM Spectrum Scale cluster information cannot be displayed. All the nodes in the cluster are down or cannot be contacted.](#)
- [HWNEP0270E The switch cannot respond to SNMP queries because of an authentication error.](#)
- [HWNEP0271E The following password decryption exception occurred: exception](#)
- [HWNEP0272E The switch cannot respond to SNMP queries because of the following exception: exception](#)
- [HWNEP0273E The following exception occurred because the OID format is incorrect: exception](#)
- [HWNEP0274E The switch cannot respond to SNMP queries because of a timeout problem.](#)
- [HWNEP0270I Retrieved the file module address file_module_address.](#)
- [HWNEP0271I No quota data was collected. Quota limits are not activated for the file systems that are associated with the IBM Spectrum Scale cluster.](#)
- [HWNEP0272I Collecting file systems that are mounted on the nodes of storage system storage_system_id.](#)
- [HWNEP0275W One or more operations failed for the CLI command that was issued for the storage system. The following error was generated: errorMsg.](#)
- [HWNEP0276E Command execution failed because sudo is not installed.](#)
- [HWNEP0277I Commands are executed through 'sudo'.](#)
- [HWNEP0278E User can not execute command through sudo.](#)
- [HWNEP0279I Collecting remote file systems for storage_system_id storage system.](#)
- [HWNEP0280I Collecting remote file systems that are mounted on the nodes of storage system storage_system_id.](#)
- [HWNEP0281E The switch is returning corrupted data.](#)
- [HWNEP0282E Zoning data cannot be collected because there is a transaction in progress on the switch](#)
- [HWNEP0283E VSAN vsan_name was not found.](#)
- [HWNEP0284E No zoning data collected from the switch.](#)
- [HWNEP0285E Cannot authenticate to the object storage using the specified user credentials.](#)
- [HWNEP0286E An object storage request failed on the GPFS cluster.](#)
- [HWNEP0287E Error when collecting Accounts information from Object Storage Service using REST protocol.](#)
- [HWNEP0288E Error when collecting Containers information from Object Storage Service using REST protocol.](#)
- [HWNEP0281I Collecting object storage accounts for storage_system_id storage system.](#)
- [HWNEP0282I Collecting object storage containers for storage_system_id storage system.](#)
- [HWNEP0289E Failed to retrieve container information because the number of containers now exceeds the maximum number of containers that can currently be collected for an account \(MAX Containers \).](#)
- [HWNEP0290E The probe failed to retrieve object storage account information from the storage system storage_system_id because the userid user does not have the required authority.](#)
- [HWNEP0291E The probe failed to retrieve object storage container information from the storage system storage_system_id because the userid user does not have the required authority.](#)
- [HWNEP0292E Cannot query the object service for information about accounts and containers as the specified user does not have admin privileges.](#)
- [HWNEP0293W The probe did not collect information about all the object accounts for the storage system storage_system_id as the userid user does not have sufficient authority on the storage system.](#)
- [HWNEP0294W An authentication error prevented the switch from responding to SNMP queries regarding the ability of the switch to perform zone control.](#)
- [HWNEP0295W A timeout prevented the switch from responding to SNMP queries regarding the ability of the switch to perform zone control.](#)
- [HWNEP0296W The switch cannot respond to SNMP queries to check the ability of the switch to perform zone control because of the following exception: exception](#)
- [HWNEP0297W The switch cannot respond to SNMP queries to check the ability of the switch to perform zone control because of the following exception: exception](#)
- [HWNEP0298I Collecting IBM Cloud Object Storage configuration.](#)
- [HWNEP0299I Collecting IBM Cloud Object Storage vaults.](#)
- [HWNEP0300I Collecting detailed IBM Cloud Object Storage status.](#)
- [HWNEP0301W The IP address ip_address for the FlashSystem storage system is not the management IP address.](#)
- [HWNEP0302I Collecting Transparent Cloud Tiering information for storage_system_id storage system.](#)
- [HWNEP0303I No Transparent Cloud Tiering configuration was detected on the IBM Spectrum Scale cluster.](#)
- [HWNEP0304E Cannot connect to IBM Cloud Object Storage.](#)
- [HWNEP0305I Collecting disk controllers for storage system storage_system_id storage system.](#)
- [HWNEP0306I Collecting disks for storage system storage_system_id storage system.](#)
- [HWNEP0307I Collecting CIFS shares for storage system storage_system_id storage system.](#)
- [HWNEP0308I Collecting NFS exports for storage system storage_system_id storage system.](#)
- [HWNEP0309I The data is being collected by the data collector: data_collector_host.](#)
- [HWNEP0310I Discovery found number storage systems.](#)
- [HWNEP0311I Probing nodes or directors for storage system name storage_system.](#)
- [HWNEP0312I Probe found number nodes or directors.](#)
- [HWNEP0313I Probing pools for storage system name storage_system.](#)
- [HWNEP0314I Probe found number pools.](#)
- [HWNEP0315I Probing disk groups for storage system name storage_system.](#)
- [HWNEP0316I Probe found number disk groups.](#)
- [HWNEP0317I Probing disks for storage system name storage_system.](#)
- [HWNEP0318I Probe found number disks.](#)
- [HWNEP0319I Probing host connections for storage system name storage_system.](#)
- [HWNEP0320I Probing ports for storage system name storage_system.](#)
- [HWNEP0321I Probing volumes for storage system name storage_system.](#)
- [HWNEP0322I Probe found number volumes. Continuing to probe volumes.](#)
- [HWNEP0323I Probe found number volumes for storage system name storage_system.](#)
- [HWNEP0324I Probing NAS nodes for storage system name storage_system.](#)
- [HWNEP0325I Probe found number NAS nodes.](#)
- [HWNEP0326I Probing file systems that are mounted on the NAS nodes of storage system name storage_system.](#)
- [HWNEP0327I Probe found number file systems.](#)

- [HWNEP0328I Probing file system exports for storage system name storage system.](#)
- [HWNEP0329W profile name version version number SMI-S Profile is not supported.](#)
- [HWNEP0330E Unable to find minimum required SMI-S profile to proceed with requested task.](#)
- [HWNEP0331I Probing copy pair relationships for storage system name storage system.](#)
- [HWNEP0332I Probe found number copy pairs.](#)
- [HWNEP1111E There is no connection for the specified device.](#)
- [HWNEP1112E No SSH server found on the device.](#)
- [HWNEP1113E Unsupported version.](#)
- [HWNEP1114E The connection to the device failed.](#)
- [HWNEP1115E Authentication failed.](#)
- [HWNEP1116E Unknown host.](#)
- [HWNEP1117E The passphrase is wrong.](#)
- [HWNEP1118E The passphrase is missing.](#)
- [HWNEP1119E Unknown error.](#)
- [HWNEP1120E ESSNI not available.](#)
- [HWNEP1121E Private key not found.](#)
- [HWNEP1122E Invalid format for the private key.](#)
- [HWNEP1123E Unable to establish a connection to the device through http port 80.](#)
- [HWNEP1124I Log collection successfully started for storage system name storage system.](#)
- [HWNEP1125E The activity requested is already in progress on storage system name storage system.](#)
- [HWNEP1126I The support log activity has started successfully storage system name storage system.](#)
- [HWNEP0112E The CLI command that was issued for the storage system failed and generated the following error: error message](#)
- [HWNEP1127I The probe failed to retrieve encryption information from the storage system storage_system_id because the userid user does not have the required authority.](#)
- [HWNEP1128E The process failed because it was unable to find the Export Tool. Expected location was loc of tool.](#)
- [HWNEP1129E The process failed because the userid or password provided failed to connect to the Export Tool.](#)
- [HWNEP1130E The process failed because the Hitachi SVP was busy and did not return data or timed out.](#)
- [HWNP5412E Performance statistics collection is not enabled.](#)
- [HWNEP1131E The process failed because the Hitachi performance interval is set to something other than 1 or 5 minutes.](#)
- [HWNEP1132W Can't collect further system information because the device returned unexpected values.](#)
- [NAD0001I Connecting to hostname using protocol protocol.](#)
- [NAD0002W Connection to hostname failed using protocol protocol: error.](#)
- [NAD0003I Connected to hostname using protocol protocol.](#)
- [NAD0005E Connection to hostname failed using protocol protocol: error message.](#)
- [NAD0006E Exception thrown for method method name: error message.](#)
- [NAD0007I Closing connection to hostname.](#)
- [NAD0008E Invalid protocol protocol passed to method name.](#)
- [NAD0010E Invalid parameter\(s\).parameter name passed to method name.](#)
- [NAD0013I Installing GUID on remote machine: hostname.](#)
- [NAD0014I GUID successfully installed on remote machine: hostname.](#)
- [NAD0018E Command on remote machine: host name failed. Error code = value executing command value.](#)
- [NAD0019E Parameter parameter passed to method is null or 0 length.](#)
- [NAD0055E Failed to connect to remote host host.](#)
- [NAD0097I Opening connection to hostname.](#)
- [NAD0180I Installing re-distributable package on .](#)
- [NAD0181I Install of re-distributable package on succeeded.](#)
- [NAD0182E Failed to install re-distributable package on .](#)
- [NAD0186I Trying to locate package TIVguid using pkginfo ...](#)
- [NAD0187I Package TIVguid is not installed.](#)
- [NAD0188I Checking TIVguid default install path : path ...](#)
- [NAD0259W Unable to determine Storage Resource Agent version on host . Fabric Discovery will not be invoked.](#)
- [NAD0145E Cannot get version information from agent on host .](#)
- [NAD0146E The connection to remote machine failed because the Remote Execution and Access component was unable to create a temporary directory on the remote machine. Remove unneeded -CSRI* directories in the remote machine's temporary directory.](#)
- [NAD0156E The server host address cannot be reached because the host name or IP address is not recognized.](#)
- [NAD0157E The server host name cannot be contacted. The server might be down, unreachable due to network problems, or the SSH credentials might be invalid.](#)
- [NAD0260I Agent is active.](#)
- [NAD0272W The connection to the Storage Resource Agent on host name was not established. Retrying using the IP address.](#)
- [NAD0274E An SSH certificate certificate name already exist.](#)
- [NAD0275E Failed to connect to remote host hostname and port. Failed to establish a secure connection.](#)
- [NAD0276E Failed to connect to remote host hostname and port. Failed to establish a secure connection because the SSL handshake failed.](#)
- [NAD0277E Failed to connect to remote host hostname and port. Failed to establish a secure connection because of an invalid SSL key.](#)
- [NAD0278E Failed to connect to remote host hostname and port. Failed to establish a secure connection because the identity of the peer could not be verified.](#)
- [NAD0279E Failed to connect to remote host hostname and port. Failed to establish a secure connection because of an SSL protocol error.](#)
- [NAD0281E The Storage Resource agent cannot be deployed because of insufficient space or other issues on the target system. The error is: error message.](#)
- [BTAVM2272W Unsupported virtual disk backing info for disk "Disk name" of hypervisor Hypervisor name, virtual machine "VM name": Virtual disk type.](#)
- [BTAVM2273W Unable to find file "File name" which is the backing device of the virtual disk "Disk name" of hypervisor Hypervisor name, virtual machine "VM name".](#)
- [BTAVM2274W Probe of hypervisor Name of the Hypervisor completed with warnings.](#)

BTADS0000I Starting Discover Process value , with Device Server RUN ID value , and Job ID value .

Explanation

The specified discovery request has been started.

Action

None.

BTADS0001I Discover Process with Device Server run ID *value* and job ID *value* is complete.

Explanation

The specified discovery request has completed.

Action

None.

BTADS0002I Starting Child Discover Process *value* with Job ID= *value* .

Explanation

The specified child discovery request has started.

Action

None.

BTADS0003I The Child Discover Process with Job ID *value* has completed with Status= *value* and Return Code= *value* .

Explanation

The specified child discovery request has completed.

Action

None.

BTADS0004W The child discovery request with job ID *job_id* completed with status *status_number* and return code *value*.

Explanation

The specified child discovery request completed with one or more warnings.

Action

For more information about the warnings, check the job log for the child discovery request. To find the child discovery request, check the job logs that have the same run ID.

BTADS0005E The child discovery request with job ID *job_id* completed with status *status_number* and return code *value*.

Explanation

The specified child discovery request completed with one or more errors.

Action

For more information about the errors, check the job log for the child discovery request. To find the child discovery request, check the job logs that have the same run ID.

BTADS0010I Invoking outband scanner value on agent value .

Explanation

The specified scan has been invoked.

Action

None.

BTADS0011I Outband scanner value on agent value completed successfully.

Explanation

The specified scan has been completed.

Action

None.

BTADS0012E Outband Scanner value on agent value failed with return code value .

Explanation

The specified scan failed.

Action

None.

BTADS0019E An outband scanner failed to capture the scan data.

Explanation

An error condition prevented the scanner from processing the outband scan data.

Action

If the problem continues, contact IBM customer support.

Related reference

-  [Getting support](#)

BTADS0020I Processing value data from agent value .

Explanation

TSANMLegacyParser is running the specified Device Server job to parse XML that represents a fabric.

Action

None.

BTADS0021W Warning encountered while parsing Fabric XML for job: RUN ID= *value* , and Job ID= *value* . *value* .

Explanation

The DESaxParser has encountered a warning while parsing XML that represents a fabric.

Action

None.

BTADS0022E Exception encountered while parsing Fabric XML for job: RUN ID= *value* , and Job ID= *value* . *value* .

Explanation

The DESaxParser has encountered an exception while parsing XML that represents a fabric.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTADS0023E Fatal error encountered while parsing Fabric XML for job: RUN ID= *value* , and Job ID= *value* . *value* .

Explanation

The DESaxParser has encountered a fatal error while parsing XML that represents a fabric.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTADS0024E Error encountered processing scanner *value* data from agent *value* . *value* .

Explanation

TSANMLegacyParser has encountered an error while parsing XML that represents a fabric.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTADS0025I Running job to discover SMI-S providers through Service Location Protocol: RUN ID= *value* , Job ID= *value* .

Explanation

SLPObjectParser is running the specified job to discover SMI-S providers through Service Location Protocol.

Action

None.

BTADS0026I Service Location Protocol has found *value* SMI-S providers.

Explanation

SLPObjectParser has found SMI-S providers through Service Location Protocol.

Action

None.

BTADS0027E Error encountered by a Service Location Protocol job: RUN ID= *value* , and Job ID= *value* . *value* .

Explanation

SLPObjectParser has encountered an error while discovering SMI-S providers through Service Location Protocol. One or more SMI-S providers may not have been discovered.

Action

Information for one or more SMI-S providers may need to be manually entered. Contact IBM customer technical support if automatic discovery of SMI-S providers through Service Location Protocol is needed.

Related reference

- [Getting support](#)

BTADS0028W The Device Server Job with RUN ID=: *value* , Job ID= *value* , Discover Request= *value* has been cancelled since it is long running.

Explanation

The specified job has been cancelled.

Action

None.

BTADS0029I Scanner *value* data from agent *value* has not changed since last scan.

Explanation

The specified scan has been invoked.

Action

None.

BTADS0030I Invoking inband Scanner *value* on agent *value* .

Explanation

The specified scan has been invoked.

Action

None.

BTADS0031I Inband Scanner *value* on Agent *value* completed successfully.

Explanation

The specified scan has been completed.

Action

None.

BTADS0032E Inband Scanner *value* failed on agent *value* with Return Code *value* .

Explanation

The specified scan failed.

Action

None.

BTADS0033E Error invoking *value* on host *value* .

Explanation

The specified scan failed.

Action

None.

BTADS0034E Fatal error encountered while persisting the data for job: RUN ID= *value* , and Job ID= *value* . *value* .

Explanation

The DESaxParser has encountered a fatal error while persisting fabric information.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTADS0035E The execution of the job failed with: *value* .

Explanation

An error occurred during the execution of a job.

Action

See the message and trace file for more information. Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTADS0036I Found SNMP Target at *value* .

Explanation

The specified SNMP agent was detected.

Action

None.

BTADS0037E Found SNMP Target at *value* but it is not persisted in the database. Will NOT perform discovery of information using the address.

Explanation

The specified SNMP agent was detected but the agent information was not found in the database.

Action

Check the message and trace files for more information.

BTADS0038I Starting scan of SNMP agents from *value* to *value* .

Explanation

Scanning the specified range of IP Addresses for SNMP agents.

Action

None.

BTADS0039I Starting probe of detected agents.

Explanation

Starting a probe against the discovered agents to collect fabric information.

Action

None.

BTADS0040I Processing of Scanner value data from Agent value completed successfully.

Explanation

The specified scan has been completed.

Action

None.

BTADS0041I Discover Process with Device Server RUN ID value and Job ID value completed successfully.

Explanation

The specified discovery request has completed successfully.

Action

None.

BTADS0042E Discover Process with Device Server RUN ID value and Job ID value failed with return code value .

Explanation

The specified discovery request failed.

Action

See the message and trace files for more information. Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTADS0043I Invoking value scanner value on target value .

Explanation

The specified scan has been invoked.

Action

None.

BTADS0044I value scanner value on target value completed successfully.

Explanation

The specified scan completed successfully.

Action

None.

BTADS0045E *value* Scanner value on target value failed with return code value .

Explanation

See the message and trace files for more information. Contact IBM customer technical support.

Action

None.

Related reference

- [Getting support](#)

BTADS0046I Processing value data from agent value .

Explanation

The specified parser is running the specified Device Server Job to parse XML from the agent.

Action

None.

BTADS0047W The value parser encountered a warning while parsing XML for job with RUN ID= value , and Job ID value . The return code from the parser job is value .

Explanation

The specified Parser has encountered a warning while parsing XML from the agent.

Action

None.

BTADS0048E The value parser encountered an exception while parsing XML from job with RUN ID= value , and Job ID= value .The return code from the parser is value .

Explanation

The Parser has encountered an exception while parsing XML from the agent.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTADS0049E The *value* parser for Device Server job with RUN ID=*value* , and Job ID= *value* failed. The return code from the parser is *value* .

Explanation

The Parser job failed

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTADS0050I Service Location Protocol has found SMI-S provider, *value* , at address *value* .

Explanation

SLPObjectParser has found SMI-S providers through Service Location Protocol.

Action

None.

BTADS0051I Service Location Protocol has found SMI-S provider, *value* , at address *value* , which requires security information to be configured.

Explanation

SLPObjectParser has found an SMI-S provider. Currently unable to communicate with the SMI-S provider because of missing or incorrect login information.

Action

Configure the login information for the SMI-S provider on the Services->Agents->CIMOMs configuration dialog under Administrative Services.

BTADS0052W Warning encountered while parsing *value* data from agent *value*. *value*.

Explanation

The DESaxParser has encountered a warning while parsing XML that represents a fabric.

Action

None.

BTADS0053E Exception encountered while parsing *value* data from agent *value*. *value*.

Explanation

The DESaxParser has encountered an exception while parsing XML that represents a fabric.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTADS0054E Fatal error encountered while parsing *value* data from agent *value*. *value*.

Explanation

The DESaxParser has encountered a fatal error while parsing XML that represents a fabric.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTADS0055E Outband Scanner *value* on agent *value* encountered the presence of a McData i10k. These devices do not report correctly via SNMP and can only be used with SMI-S provider.

Explanation

McData i10k devices do not work properly with SNMP agents. They should only be used with SMI-S providers. Remove this SNMP agent.

Action

Remove this device as an Out of band Fabric agent and add it as an SMI-S provider.

BTADS0056E Errors in Topology XML generator.

Explanation

Some of the data generated by the scanner appears to have errors.

Action

Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTADS0057E Errors occurred while resolving InterconnectElement and Port relationship.

Explanation

Some of the data returned from the scanner could not be resolved.

Action

Get the SANQueryEngine service trace information from the trace log, and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTADS0058E Errors in creating an entity.

Explanation

Errors occurred while attempting to create an entity from the results of a scan.

Action

Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTADS0059E The outband agent target address *IP address* is not a Cisco device or is invalid.

Explanation

The Cisco scan cannot be run on the outband agent target address. Either the device is not a Cisco device or the IP address is not valid.

Action

Verify that the connection information for the device is valid. Ensure that the network and the device are up and available. Try the action again.

BTADS0060E Outband Scanner *value* is not responding.

Explanation

The specified scan failed has failed due to a timeout.

Action

Please check network connections to the switch and SNMP settings on the switch.

BTADS0062E Encountered SQL error *value* while persisting some data.

Explanation

Some discovered data was not stored in the database due to an SQL error.

Action

Please check the device server trace log for more detail. If the problem continues, contact IBM support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTADS0063E The execution of the PM BSP invocation failed with:
value .

Explanation

An error occurred during the PM BSP invocation.

Action

See the message and trace file for more information. Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTADS0063W The performance data collection for the current device
is not enabled.

Explanation

The performance collection service is disabled. Please check the external provider documentation.

Action

None.

BTADS0064I Starting scan of Storage Subsystems from *value* to *value*
.

Explanation

Scanning the specified range of IP Addresses for Storage Subsystems.

Action

None.

BTADS0065I Outband and inband agents for fabric(s) specified in
probe are *value*

Explanation

Determining outband and inband agents that can be used to probe fabric(s)

Action

None.

BTADS0066I Could not find scanners for agent *value*

Explanation

Unable to find scanners for agent and so will not be able to invoke the scanners for this particular agent during the probe

Action

None.

BTADS0067I Agent value is configured for no SAN calls and so no scanners will be invoked for this particular agent

Explanation

Determining outband and inband agents that can be used to probe fabric(s)

Action

None.

BTADS0068I Could not retrieve connection information for agent value. Will not be able to invoke scanners for this particular agent

Explanation

Unable to invoke scanners for a agent since connection information for the agent could not be obtained

Action

None.

BTADS0069I Added inband scanner job with id value discover request value for agent value.

Explanation

Added a job to invoke scanner for a particular agent to probe a fabric

Action

None.

BTADS0070I Agent value has not discovered any fabrics and will not be used during the probe.

Explanation

This particular agent has not discovered a fabric previously and will not be used during this probe run

Action

None.

BTADS0071I Invoked inband Scanner value on agent value .

Explanation

The specified scan has been invoked successfully.

Action

None.

BTADS0072I Successfully received response from agent for job value with request id value .

Explanation

Response received from agent successfully. Server will process the response.

Action

None.

BTADS0073E Received error response from agent for job value with request id value. Return code is value.

Explanation

Error response received from agent for specified job or there was a error encountered while receiving job. Job will be considered failed.

Action

None.

BTADS0074E IP Scan Discovery was canceled due to a hung socket/thread detected. Partial result of the scan will be persisted.

Explanation

IP Scan Discovery was canceled since a hung thread or socket was detected. Available results will be persisted in the database.

Action

None.

BTADS0075E IP Scan Discovery was canceled due to a hung socket/thread detected.

Explanation

IP Scan Discovery was canceled since a hung thread or socket was detected.

Action

None.

BTADS0076I IP Scan Discovery has started for DS, XIV, and IBM SONAS subsystems .

Explanation

IP Scan Discovery for DS, XIV, and IBM SONAS subsystems has started

Action

None.

BTADS0077I Scanning value out of value IP addresses.

Explanation

Informational message on number of IP being scanned.

Action

None.

BTADS0078I IP Scan Discovery has started for SVC subsystems.

Explanation

IP Scan Discovery for SVC subsystems has started

Action

None.

BTADS0079I IP Scan Discovery for DS and XIV was done.

Explanation

IP Scan Discovery for DS, XIV, and IBM SONAS was done

Action

None.

BTADS0080I IP Scan Discovery for SVC was done

Explanation

IP Scan Discovery for SVC was done.

Action

None.

BTADS0081I Inband Scanner value for agent address value is not required for probing switches and will not be used.

Explanation

The information that can be collected using the specified scanner will be collected by another scanner or agent instead.

Action

None.

BTADS0082W A first run of a switch probe failed. Additional agents will be used.

Explanation

A probe that used a subset of agents to probe switches did not complete. The probe will now use previously unassigned agents. Depending on the types of agents that are available, all or a subset of information will be collected about the switches.

Action

Check the trace log for more detail. Determine whether the agent or agents that failed are currently down for a known reason. If the problem continues, contact IBM Software Support.

BTADS0083I The available agents provide a subset of possible features for the probed switch: *value*

Explanation

Some switches require more than one type of agent (such as a Storage Resource agent, SNMP agent, and SMI agent) to collect all the available information for the switch.

Action

To determine if any additional agent types are required for the switch and fabric functions that you need, go IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93> and view the Reference section. If you have all the required agents deployed for the switch, ensure that are configured and running properly.

BTADS0084I There are no limitations for probing switch *value* based on the mix of agents that are configured.

Explanation

With the current mix of configured agents, there are no known limitations on the information that can be collected from a probe of the specified switch.

Action

No action is required.

BTADS0085W A problem was encountered when agent assignments were being determined for the probe. All available agents will be used to collect information about the switch.

Explanation

Typically, agent assignment is used to reduce the number of agents that are used by a probe. Because an exception was encountered while determining agent assignments, all the available agents will be used.

Action

Check the device server message and trace log for more information. If the problem continues, contact IBM support.

BTADS0086I The following storage systems were discovered *value*

Explanation

This message is for informational purposes only.

Action

No action is required.

BTADS0087I IP Scan Discovery did not find any DS8000, SVC, XIV, and IBM SONAS storage systems in the given IP range.

Explanation

This message is for informational purposes only.

Action

No action is required.

BTADS0088I IP Scan Discovery finished.

Explanation

This message is for informational purposes only.

Action

No action is required.

BTADS0089E The Device server is not registered with agent manager. Scanners cannot be used for agent value. value.

Explanation

The Device server must be registered with the agent manager.

Action

Configure the Device server to register with the agent manager. Check the Device server message and trace log for more information. If the problem continues, contact IBM support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTADS0090E There are no agents currently available to probe switch value.

Explanation

There are no agents configured to probe the specified switch or the agents configured are not operational.

Action

Configure the agents to probe switches. If agents are already configured, check if they are operational. Check the trace log for more detail. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTADS0091I Inband Scanner value for agent address value is currently not running and will not be used.

Explanation

The specified scanner is currently not running and will not be used.

Action

Check the status of specified agent and determine the reason that it is not running. Check the device server message and trace log for more detail. Check if an alternate agent can be used for the probe. If the problem continues, contact IBM support.

BTADS0092I Inband Scanner value for agent address value is currently disabled from performing fabric functions and will not be used.

Explanation

The specified scanner must be enabled if you want to perform fabric functions.

Action

Check the reason for disabling the fabric functions for the specified agent. Check if an alternate agent can be used for the probe.

BTADS0093I Inband Scanner value for agent address value is currently not reachable and will not be used.

Explanation

The specified scanner must be reachable to use it for a probe.

Action

Check if the host on which the agent is running is reachable from the server and that the server is reachable from the agent. Check the trace log for more detail. Check if an alternate agent can be used to probe switches. If the problem continues, contact IBM Software Support.

BTADS0094W The probe for switch value has some limitations.

Explanation

There are some limitations for the information collected by a probe of the switch based on the mix of agents that are configured and operational.

Action

Check if there are additional agents that can be configured for the probe. Check the status of the configured agents. Check the device server message and trace log for more detail. If the problem continues, contact IBM support.

BTADS0095W For switch value some information will not be collected.

Explanation

There are some limitations for the information collected by a probe of the fabric based on the mix of agents that are configured and operational.

Action

Check if there are additional agents that can be configured for the probe. Check the status of the configured agents. Check the device server message and trace log for more detail. If the problem continues, contact IBM support.

BTADS0096I The probe limitation can be overcome by configuring an SMI agent to manage fabric value.

Explanation

An SMI agent is required to manage the fabric. Note that there might be more than one type of agent that can manage the fabric.

Action

Configure an SMI agent to manage the specified fabric. For more information about the types of agents that are available to manage a fabric, go IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93> and view the Reference section.

BTADS0097I The probe limitation can be overcome by configuring SNMP agents to manage switches in fabric value.

Explanation

SNMP agents are required to manage the specified fabric. Note that there might be more than one type of agent that can manage the fabric.

Action

Configure an SNMP agent to manage the specified fabric. For more information about the types of agents that are available to manage a fabric, go IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93> and view the Reference section.

BTADS0098I The probe limitation can be overcome by configuring a Storage Resource agent to manage fabric value.

Explanation

A Storage Resource agent is required to manage the specified fabric. Note that there might be more than one type of agent that can manage the fabric.

Action

Configure a Storage Resource agent to manage the specified fabric. For more information about the types of agents that are available to manage a fabric, go IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93> and view the Reference section.

BTADS0099W The following WWN is not recognized as belonging to a known vendor: value.

Explanation

A unique identifier contained in the WWN was not found and associated with a known vendor.

Action

Contact IBM support for more information about associating the identifier in the WWN with a corresponding vendor.

BTADS0100W Invalid relationships between switches and fabrics were identified. If possible, these relationships will be fixed automatically for the following switches: value.

Explanation

Some switches were identified as belonging to more than one fabric. This issue might occur if two fabric data sources report different principal switches for the same fabric. The firmware on the switches might be the root cause of this problem. One of the principal switches is chosen to be displayed as the fabric WWN and the other switch WWN is not shown as a separate fabric.

Action

Ensure that all switches in your fabric are running the same firmware level. If the problem continues, upgrade the switches to the latest firmware level. If the problem still occurs, contact IBM customer support.

BTADS0101W The discover process that has the Device server run ID value and job ID value completed with one or more warnings.

Explanation

The specified discovery process completed with one or more warnings.

Action

See the probe logs for this run ID to view the warning messages.

BTADS0102E The probe with the run ID value completed with errors.

Explanation

Configuration, capacity, and status information couldn't be collected about the resource.

Action

Complete one or more of the following actions:

- Check the probe log for related error messages that might help determine the cause of the problem.
- Verify that the resource that is being monitored is up and available.
- Verify that the local area network is available and that the firewall is not preventing network access to product services and agents. Contact your firewall administrator to ensure that the required ports are open.
- Check the status of the product servers and database repository on the Home > System Management page.
- On Windows, verify that the related database services are active.
- Check for error messages in the log files for the servers.
- Try the probe again. If the problem persists, contact IBM Software Support.

Related reference

- [Ports used by IBM Spectrum Control](#)
- [Getting support](#)

BTADS0103E No data source is available to probe switch *switch_name*.

Explanation

No data source is configured to probe the specified switch or the data source that is configured is not operational.

Action

Configure a data source to probe the switch. If a data source is already configured, check whether it is operational. Also, check the message and trace log on the Device server for more detail. If you cannot resolve the problem, contact IBM Software Support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTADS0104E A timeout occurred while processing the request. Try the request again.

Explanation

The request could not be processed in the time allowed.

Action

Try the request again.

BTADS0105E A response from the data collector was not received within the specified time.

Explanation

The data collector did not respond to the server in the specified time. The data collector might not be running or it might not be able to connect to the server.

Action

Verify that the data collector is running and that it can connect to the server. Then try the request again.

BTADS0106E The requested action on agent *agent_name* did not complete because the data collector stopped or is not responding. The request failed with error code *error_code*.

Explanation

The specified scan did not complete.

Action

Verify that the data collector is running and that it can connect to the server. Then try the request again.

BTADS0107W Outband Scanner *outband_scanner_name* on agent *agent_name* failed because of another transaction in progress on the switch.

Explanation

The switch is busy handling another transaction. To avoid this error, wait for the transaction to be done and try to run the probe again.

Action

Check the trace file for more details. Try the operation again.

BTADS0108E Outband Scanner *outband_scanner_name* on agent *agent_name* failed because unexpected data was returned by the switch. Check the trace file for more details.

Explanation

The switch is returning unexpected data. To avoid this error, review recent configuration changes and try to run the probe again.

Action

Check the trace file for more details. Try the operation again.

BTADS0109I Outband Scanner *outband_scanner_name* on agent *agent_name* did not collect zoning data.

Explanation

The switch did not return zoning data. There may be issues with the switch configuration or zoning configuration.

Action

Review recent configuration changes. Try the operation again.

BTADS0110I Outband Scanner *outband_scanner_name* on agent *agent_name* did not pass write capabilities check.

Explanation

The IBM Spectrum Control server could not verify write capabilities either due to an incorrect write community string or an authentication failure. Without correct credentials, zone control functions will not work.

Action

Check and if necessary correct either the write community string or the authentication parameters.

BTADS0111E The probe was unable to collect some details of the switch.

Explanation

A problem was encountered during the probe of the switch. The switch properties might be reported incorrectly until a later probe is completed.

Action

Try the following actions to resolve the problem:

- Check the logs for an indication of the error or exception, and resolve the problem if possible. For information about probe logs, see the Viewing probe logs topic in the IBM Spectrum Control Knowledge Center.
- Try the probe again at a later time. For instructions on how to start a probe, see the Starting probes topic in the IBM Spectrum Control Knowledge Center.
- If the problem continues, contact IBM customer support.

Related reference

-  [Getting support](#)

BTADS0112E Error encountered while persisting some data. *value*

Explanation

Some discovered data was not stored in the database due to an error.

Action

See the message and trace file for more information. Contact IBM customer for technical support.

Related reference

-  [Getting support](#)

BTADS0113E Error encountered while processing a probe job. *value*

Explanation

An error encountered while processing a probe job.

Action

See the message and trace file for more information. Contact IBM customer for technical support.

Related reference

- [Getting support](#)

BTADS0114E The information cannot be saved to the database repository.

Explanation

A problem occurred during a database operation.

Action

Try the following actions:

- Verify that the local area network is available.
- Check the status of the database repository on the Home > System Management page.
- On Windows, verify that the related database services are active.
- If you need more information, go to the IBM Knowledge Center and view the information in the "Troubleshooting and problem determination" section.
- If you cannot resolve the problem, contact IBM Software Support.

Related reference

- [IBM Spectrum Control documentation](#)
- [Getting support](#)

BTADS0115E The probe failed when collecting information about the resource. The data collector returned the following error status: *value*.

Explanation

Configuration, capacity, and status information couldn't be collected about the resource.

Action

Check the probe log for related error messages that might help determine the cause of the problem.

BTAFM0000I Operation *op_name* processed successfully.

Explanation

A Device server API has been executed successfully.

Action

No action is required.

BTAFM0100I Initializing Collection.

Explanation

Discovery or probe is being initialized.

Action

No action is required.

BTAFM0110I Querying the SMI-S provider.

Explanation

An SMI-S provider query is in process.

Action

No action is required.

BTAFM0113I Collecting for *db_table*, *current_obj* of *num_objs*.

Explanation

This is a discovery or probe status message indicating the Fabric objects for which data is being collected, and the expected number of objects to build.

Action

No action is required.

BTAFM0114I Probing data for switch *switch_name*.

Explanation

The attributes are being gathered for a particular switch.

Action

No action is required.

BTAFM0115I Probing data for port *port_name*.

Explanation

The attributes are being gathered for a particular port.

Action

No action is required.

BTAFM0150I Storing Information.

Explanation

The discovery or probe is currently storing collected data.

Action

No action is required.

BTAFM0151I The *db_table* of *current_obj* *num_objs* stored.

Explanation

The discovery or probe is currently storing information. The progress of the discovery or probe is contained in the message.

Action

No action is required.

BTAFM0200I Traversing fabric topology.

Explanation

The discovery or probe is currently traversing CIM classes that are related to switches and ports.

Action

No action is required.

BTAFM0500I The IBM Spectrum Control Device Server service has started successfully.

Explanation

The startup of the IBM Spectrum Control Device Server service has completed without error.

Action

No action is required.

BTAFM0501I The IBM Spectrum Control Device Server service was shut down successfully.

Explanation

The shutdown of the IBM Spectrum Control Device Server service has completed without error.

Action

No action is required.

BTAFM0502I The IBM Spectrum Control Device Server service provides methods to collect, report and configure the fabric hardware.

Explanation

This is a description of the Device Server service component of IBM Spectrum Control.

Action

No action is required.

BTAFM0505I The delete missing function has started.

Explanation

The delete missing function has started to remove objects in the 'missing' state from the database.

Action

No action is required.

BTAFM0506I The delete missing method was processed in *milliseconds* milliseconds.

Explanation

The delete missing function finished successfully and objects in the 'missing' state were removed from the database.

Action

No action is required.

BTAFM0723W No blades were discovered for the slot *slot*.

Explanation

The information for the blade was not added to the database.

Action

Check the trace log and Device server log for error messages that might help determine the problem. For information about the location of log files, go to the documentation at <https://www.ibm.com/docs/en/spectrum-control/>.

BTAFM2000W Operation *op_Name* partially processed.

Explanation

A request has been partially executed. Some operations as part of this request have failed.

Action

Check the logs for an indication of an error or exception and contact IBM customer support.

BTAFM2501W Unable to shut down Device Server Service smoothly.

Explanation

An exception occurred while shutting down one or more components.

Action

No further action needed since the service will be forced to shut down.

BTAFM4000E Operation *op_Name* failed.

Explanation

A request has failed. None of the operations that are part of this request were processed successfully.

Action

Check the logs for an indication of an error or exception and contact IBM customer support.

Related reference

- [Getting support](#)

BTAFM4001E An internal error occurred.

Explanation

An internal error occurred during execution.

Action

Check the logs for an indication of an error or exception and contact IBM customer support.

Related reference

- [Getting support](#)

BTAFM4002E Could not get requested information due to an internal error - *errorMessage*

Explanation

An error occurred during the attempt to retrieve information from the database.

Action

Check the health of the database. Then, rerun the discovery and probe process, and then run a data collection task.

BTAFM4100E Mandatory parameter *parameter_Name* is missing.

Explanation

A service has been called without a required parameter.

Action

Check the mandatory parameters and retry the operation.

BTAFM4101E Invalid parameter *parameter_name*.

Explanation

A service has been called with an invalid parameter.

Action

Check the valid parameters and retry the operation.

BTAFM4103E Entity *entity_name* was not found.

Explanation

The database entity that was specified was not found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4104E Attribute *attribute_name* was not found.

Explanation

The database column specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4105E Computer *computer_name* was not found.

Explanation

No computer system was found that corresponds to the key that was passed in as an input parameter. The computer system specified as a parameter can not be found.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4106E Fabric *fabric_name* was not found.

Explanation

The fabric specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4107E Switch *switch_name* was not found.

Explanation

The switch specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4108E Port *port_name* was not found.

Explanation

The port specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4109E Zone set *zoneset_name* was not found.

Explanation

The zone set specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4110E Zone *zone_name* was not found.

Explanation

The zone specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4111E Zone alias *zone_alias_name* was not found.

Explanation

The zone alias specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4112E Zone member *zone_member_name* was not found.

Explanation

The zone member specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4113E Subsystem *subsystem_name* was not found.

Explanation

The subsystem specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4114E Host Bus Adapter *HBA_name* was not found.

Explanation

The Host Bus Adapter specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4115E Node *node_name* was not found.

Explanation

The node specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4116E Link from port *from_port_name* to port *to_port_name* was not found.

Explanation

The link specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4117E Hub *hub_name* was not found.

Explanation

The hub port specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4118E Router *router_name* was not found.

Explanation

The router specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4119E Bridge *bridge_name* was not found.

Explanation

The bridge specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4120E LUN *LUN_name* was not found.

Explanation

The LUN specified as a parameter can not be found in the database.

Action

Rerun the discovery and probe process, and then run a data collection task.

BTAFM4140E Agent *Agent_name* was not found.

Explanation

The agent specified as a parameter can not be found in the database.

Action

Open the agent configuration panel and check on the status of the agents.

BTAFM4141E Scanner *scanner_name* on agent *agent_name* was not found.

Explanation

The scanner specified as a parameter can not be found in the database.

Action

Open the agent configuration panel and check on the status of the agents.

BTAFM4142W Agent *agent_name* was ignored because the switch *switch_name* was probed by agent *agent1_name*.

Explanation

The agent specified was ignored because another, more appropriate one, was used for the same WWN.

Action

No action is required.

BTAFM4150E Indexed properties *property_name* don't match.

Explanation

Indexed input parameters different sizes.

Action

Make sure that all the indexed input parameters are the same size.

BTAFM4180E Agent to gather sensor and event data is not available for the switch *switch_name*.

Explanation

The switch event and sensor data is only available through SNMP agents.

Action

Make sure you have configured the correct IP address for the switch for which you want to obtain sensor and event data.

BTAFM4200E Credentials not found.

Explanation

A problem occurred while accessing the user credentials.

Action

Make sure the database is running and that the IBM Spectrum Control setup completed successfully.

BTAFM4300E The connection to the SMI agent for switch *switch_name* could not be made.

Explanation

The creation of the SMI agent client failed.

Action

Ensure that the SMI agent is running and that the correct user ID and password have been set for this SMI agent in the user interface panels.

BTAFM4301E The invocation of CIM method *method_name* failed on SMI-S provider *SMI-S provider name*. The return code is *return_code*.

Explanation

The CIM method that was invoked on the specified SMI-S provider failed.

Action

Check the health of the SMI-S provider. Check the trace log for an indication of an error or exception and contact IBM customer support.

Related reference

- [Getting support](#)

BTAFM4302E The invocation of CIM method *method_name* failed on SMI-S provider *SMI-S provider name* with the following exception text: *exception_text*.

Explanation

The CIM method that was invoked on the specified SMI-S provider failed.

Action

Check the health of the SMI-S provider. Check the trace log for an indication of an error or exception and contact IBM customer support.

Related reference

- [Getting support](#)

BTAFM4303E Received unexpected values from SMI-S provider *SMI-S provider name* .

Explanation

The SMI-S provider returned values that are unexpected and might indicate that the SMI-S provider is not working correctly.

Action

Using the CIMBrowser, make sure the SMI-S provider has the correct information. If not, fix the problem on the SMI-S provider. Otherwise, check the logs for an indication of an error or exception and contact IBM customer support.

Related reference

- [Getting support](#)

BTAFM4304E SMI agent *SMI agent name* can not contact switch *switch_name*.

Explanation

The SMI agent can not contact the switch that it manages.

Action

Check the SMI agent logs to determine why it cannot contact the switch.

BTA FM4305E The CIM method *method_name* is not supported on the switch *switch_name*.

Explanation

The switch does not support the specified function.

Action

No action is required.

BTA FM4306E Could not create connection to SMI-S provider *SMI-S provider name* . Reason: *reason*.

Explanation

A connection to the SMI-S provider could not be established.

Action

Check the health of the SMI-S provider. Check the trace log for an indication of an error or exception and contact IBM customer support.

Related reference

- [Getting support](#)

BTA FM4307E The username *user_name* or password is wrong on SMI-S provider *SMI-S provider name*.

Explanation

The username or password specified for the SMI-S provider are not valid.

Action

Specify the correct username/password combination for the SMI-S provider.

BTA FM4308I Could not create connection to SMI-S provider *SMI-S provider name* . Reason: *reason*. An alternate SMI-S provider will be used.

Explanation

A connection to the SMI-S provider could not be established. An alternate SMI-S provider has been identified to be used to manage the fabric or switch.

Action

Ensure that the server and the SMI-S provider service are running properly.

Check the dmSvcTrace.log file for an indication of the error or exception, and for the URL of the problematic SMI-S provider. This file is located in installation_directory\IBM\TPC\device\log. Contact IBM customer support.

BTA FM4501E No agent is available to configure the zoning on the fabric with ID *fabric_name*.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric. An agent to change the zone configuration was not found.

Action

Check the health of all agents for the fabric. Check that the user id and password are valid.

BTA FM4502E The fabric with ID *fabric_name* is currently locked by another client of IBM Spectrum Control.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric. The fabric is currently reserved by another client of IBM Spectrum Control. This fabric lock is used to prevent other clients of this instance of IBM Spectrum Control from attempting concurrent zone configuration changes to the same fabric.

Action

Try the zone configuration again after the other client of IBM Spectrum Control has finished.

BTA FM4503E A token for fabric *fabric_name* has expired for client *client_name*.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric. The token specified has expired.

Action

The client must call reserve again to get a new token.

BTA FM4504E The transaction for fabric *fabric_name* has expired.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric. The transaction has expired.

Action

Start a new transaction to continue the control operations.

BTA FM4505E Another transaction is in progress for fabric *fabric_name*.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric. Another transaction is already in progress.

Action

Wait until the other client finishes before starting a new transaction.

BTA FM4506E Zone set *zoneset_name* already exists.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric because a zone set already exists with the specified name.

Action

Retry the operation with a unique zone set name.

BTAFM4507E Zone *zone_name* already exists.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric because a zone already exists with the specified name.

Action

Retry the operation with a unique zone name.

BTAFM4508E Zone alias *zone_alias_name* already exists.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric because a zone alias already exists with the specified name.

Action

Retry the operation with a unique zone alias name.

BTAFM4509E Zone member *zone_member_name* already exists.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric because a zone member already exists with the specified name.

Action

Retry the operation with a unique zone member name.

BTAFM4510E Another job is in progress for fabric *fabric_name*.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric. Another job is already in progress.

Action

Wait until the other job finishes before starting a new transaction.

BTAFM4550E The Device Server encountered an error accessing the database.

Explanation

The Device Server service cannot access data stored in the database.

Action

Check the status of the database. Also, check the logs for an indication of an error or exception and contact IBM customer support.

Related reference

- [Getting support](#)

BTAFM4600E Unable to start the Device Server service.

Explanation

An exception occurred while starting up one or more components.

Action

If this is a new installation, make sure the install procedures have been followed. Also, check the trace log for an indication of an error or exception and contact IBM customer support.

Related reference

- [Getting support](#)

BTAFM5000E Step failed after collecting *Count of collected entities* entities for switch *switch* where entities exist. Continuing with next step.

Explanation

During a probe job one of the collection steps has failed. This entity type is probably not entirely collected. The probe will continue to execute the subsequent collection steps.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTAFM5001E No set of fabrics or switches was defined for this probe.

Explanation

No set of fabric WWNs or switch WWNs was passed as an input argument to the probe.

Action

If this error occurs during a scheduled probe or a probe that starts immediately, reschedule the probe.

If this error occurs in the message logs for a probe that was run by an event, you can ignore the error if no other problems are observed.

BTAFM5002E The SMI agents *SMIURL* returned an error or can no longer contact the switches.

Explanation

When a fabric uses a data source from an SMI agent, the fabric probe uses the data source to try to rediscover the fabrics and switches. A problem was encountered during this step. The SMI agents either returned an error, or the SMI agents can no longer contact the switches.

Action

Check that the SMI agent service is running. Check the SMI agent logs and UI for problems, for example, problems with contacting the switches. Check for any general problems on the SMI agent system. For example, check the memory available on the system, and check that the disk drive space is sufficient. If you cannot resolve the problem, contact IBM Software Support.

Related reference

-  [Getting support](#)

BTAFM5003E Requests to an SMI agent did not correctly collect a set of switches for fabric *fabric identity*.

Explanation

The probe of an SMI agent did not collect a set of switches for the fabric. This may indicate that the SMI agent was not able to contact the switches during the probe.

Action

Check that the SMI agent service is running. Check the SMI agent logs and UI for problems, for example, problems with contacting the switches. Check for any general problems on the SMI agent system. For example, check the memory available on the system, and check that the disk drive space is sufficient. If you cannot resolve the problem, contact IBM Software Support.

Related reference

-  [Getting support](#)

BTAFM5004E No switch retrieved from the SMI agent for fabric *fabric identity*.

Explanation

During a probe job associator call from Fabric to Switch returns nothing.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

-  [Getting support](#)

BTAFM5005E No switch found for fabric *fabric identity*.

Explanation

During a probe job no switch found for fabric.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

-  [Getting support](#)

BTAFM5006E No switch retrieved from database.

Explanation

During a probe job cannot retrieve switch from database.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTAFM5007E Failed to get CIM entity for fabric *fabric_name*.

Explanation

Failed to get fabric CIM entity.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTAFM5008E Failed to get CIM entity for switch *switch_name*.

Explanation

Failed to get switch CIM entity.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTAFM5009E Failed to enumerate CIM entity *Entity class name*.

Explanation

Failed to enumerate CIM entity.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTAFM5010E SMI-S provider is not available.

Explanation

SMI-S provider is not available.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTAFM5011E Failed to get blade for switch *Switch name*.

Explanation

Failed to get blade for switch.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTAFM5012E Failed to get physicalpackage for blade with slot number *Blade slot name*.

Explanation

Failed to get blade physicalpackage.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTAFM5013E Blade serial number is NULL.

Explanation

Blade serial number is NULL.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTAFM5014E Step failed after collecting *Count of collected entities* entities for fabric *fabric where entities exist*. Continuing with next step.

Explanation

During a probe job one of the collection steps has failed. This entity type is probably not entirely collected. The probe will continue to execute the subsequent collection steps.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTA FM5015E Data source could not be retrieved from the IBM Spectrum Control database for fabric *fabric* where data source *exists*.

Explanation

Data source could not be retrieved from the IBM Spectrum Control database.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTA FM5016E The selected data source could not be contacted for fabric *fabric* where data source *exists*.

Explanation

The inactive data source could not be contacted.

Action

Make sure the selected data source is up running. If not, select another data source and re-run the probe.

BTA FM5017E Failed to get fabric for switch *Switch name*.

Explanation

Failed to get fabric for switch.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTA FM5018E Failed to get CIM entity for port *port_name*.

Explanation

Failed to get port CIM entity.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTAFM5019E Failed to get switch for port *port_name*.

Explanation

Failed to get switch for port.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

-  [Getting support](#)

BTAFM5020E Failed to get blade for port *port_name*.

Explanation

Failed to get blade for port.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

-  [Getting support](#)

BTAFM5021E Failed to get CIM entity for blade *blade_name*.

Explanation

Failed to get blade CIM entity.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

-  [Getting support](#)

BTAFM5022E Failed to get switch for blade *blade_name*.

Explanation

Failed to get blade switch.

Action

Look for prior error messages in this log and review the traceFabric.log file. If the problem persists, contact IBM Support.

Related reference

-  [Getting support](#)

BTAFM5023E Failed to discover Fabric and Switch.

Explanation

Failed to discover Fabric and Switch.

Action

Look for prior error messages in this log and review the device discover log file. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

BTAFM5024E The data source for switch *switch_name* was not retrieved from the database repository.

Explanation

The data source for the switch was not retrieved from the database repository. This error might occur if the database repository is unavailable, if an internal error such as an SQL exception occurred, or if the data source for the switch was removed.

Action

Check the log for the probe for previous error messages that might help determine the cause of the problem. Review the trace log for the fabric.

If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

BTAFM0600I *Count of collected entities* blades collected for switch *switch* where *entities* exist.

Explanation

A probe job has collected the given number of blades so far.

Action

No action is required.

BTAFM0601I Starting collection of switch blades for switch *switch identifier*.

Explanation

A probe job has begun to collect entities from the given switch.

Action

No action is required.

BTAFM0602I Collection of switch blades completed. *Count of collected entities* entities collected in total for switch *switch identifier*.

Explanation

A probe job has collected the given number of entities.

Action

No action is required.

BTAFM0603I Starting collection of switch fcports for switch *switch identifier*.

Explanation

A probe job has begun to collect entities from the given switch.

Action

No action is required.

BTAFM0604I Collection of switch fcports completed. *count of collected entities* entities collected in total for switch *switch identifier*.

Explanation

A probe job has collected the given number of entities.

Action

No action is required.

BTAFM0605I Start probing *switch entities* switches.

Explanation

A probe job starts to probe the given number of switch entities.

Action

No action is required.

BTAFM0606I Start topology probing for fabric *fabric entity*.

Explanation

A topology probe job starts to probe the given fabric entity.

Action

No action is required.

BTAFM0609I *Count of entities* fcports collected for switch *switch where entities exist*.

Explanation

A probe job has collected the given number of fcports so far.

Action

No action is required.

BTAFM0614I The probe task is to probe topology and zone. The probe algorithm is CIM association.

Explanation

The fabric topology and zone probe algorithm is to use a series of associator requests to the SMI-S provider to collect fabric inventory.

Action

No action is required.

BTAFM0616I The probe policy involves discovering segmented or merged fabrics.

Explanation

The probe policy will discover segmented or merged fabrics.

Action

No action is required.

BTAFM0617I The probe policy doesn't involve discovering segmented or merged fabrics.

Explanation

The probe policy will not discover segmented or merged fabrics.

Action

No action is required.

BTAFM0618I The probe task is to probe topology. The probe algorithm is CIM association.

Explanation

The fabric topology probe algorithm is to use a series of associator requests to the SMI-S provider to collect fabric inventory.

Action

No action is required.

BTAFM0620I Start zone probing for fabric *fabric entity*.

Explanation

A zone probe job starts to probe the given fabric entity.

Action

No action is required.

BTA FM0621I Starting collection of zone set for switch *switch entity*.

Explanation

A probe job has begun to collect entities from the given switch.

Action

No action is required.

BTA FM0622I Starting collection of zone for switch *switch entity*.

Explanation

A probe job has begun to collect entities from the given switch.

Action

No action is required.

BTA FM0623I Starting collection of zone alias for switch *switch entity*.

Explanation

A probe job has begun to collect entities from the given switch.

Action

No action is required.

BTA FM0624I Starting collection of zone member from zone alias for switch *switch entity*.

Explanation

A probe job has begun to collect entities from the given switch.

Action

No action is required.

BTA FM0625I Starting collection of zone member and zone alias from zone for switch *switch entity*.

Explanation

A probe job has begun to collect entities from the given switch.

Action

No action is required.

BTA FM0626I Starting collection of zone member from zone for switch *switch entity*.

Explanation

A probe job has begun to collect entities from the given switch.

Action

No action is required.

BTA FM0627I Starting collection of zone set for fabric *fabric entity*.

Explanation

A probe job has begun to collect entities from the given fabric.

Action

No action is required.

BTA FM0628I *Count of collected entities* zone sets collected.

Explanation

A probe job has collected the given number of zone sets so far.

Action

No action is required.

BTA FM0629I Collection of zone set completed. *Count of collected entities* entities collected in total for fabric *fabric entity*.

Explanation

A probe job has collected the given number of entities.

Action

No action is required.

BTA FM0630I Starting collection of zone for fabric *fabric entity*.

Explanation

A probe job has begun to collect entities from the given fabric.

Action

No action is required.

BTA FM0631I *Count of collected entities* zones collected.

Explanation

A probe job has collected the given number of zone sets so far.

Action

No action is required.

BTAFM0632I Collection of zone completed. *Count of collected entities entities collected in total for fabric fabric entity.*

Explanation

A probe job has collected the given number of entities.

Action

No action is required.

BTAFM0633I Starting collection of zone alias for fabric *fabric entity.*

Explanation

A probe job has begun to collect entities from the given fabric.

Action

No action is required.

BTAFM0634I *Count of collected entities zone aliases collected.*

Explanation

A probe job has collected the given number of zone sets so far.

Action

No action is required.

BTAFM0635I Collection of zone alias completed. *Count of collected entities entities collected in total for fabric fabric entity.*

Explanation

A probe job has collected the given number of entities.

Action

No action is required.

BTAFM0636I Starting collection of zone member from zone alias for fabric *fabric entity.*

Explanation

A probe job has begun to collect entities from the given fabric.

Action

No action is required.

BTAFM0637I Starting collection of zone member and zone alias from zone for fabric *fabric entity*.

Explanation

A probe job has begun to collect entities from the given fabric.

Action

No action is required.

BTAFM0638I Starting collection of zone member from zone for fabric *fabric entity*.

Explanation

A probe job has begun to collect entities from the given fabric.

Action

No action is required.

BTAFM0639I Collection of zone member completed. *Count of collected entities* entities collected in total for fabric *fabric entity*.

Explanation

A probe job has collected the given number of entities.

Action

No action is required.

BTAFM0640I Zone probe will discover both active and inactive zone definitions at selected data source *datasource name for zone probe*.

Explanation

Zone probe job discovers both active and inactive zone definitions on a selected data source.

Action

No action is required.

BTAFM0641I Zone probe will discover only active zone sets at data source *datasource name for zone probe*.

Explanation

Zone probe job discovers only active zone sets on an alternative data source.

Action

No action is required.

BTAFM0654I The port is not switch port.

Explanation

This is not switch port.

Action

No action is required.

BTAFM0655I The switch profile doesn't support this switch *switch_name*. No further process to probe this switch.

Explanation

This is not a switch registered profile supported switch. No further process to probe this switch.

Action

No action is required.

BTAFM0656I Start enumerating entity of association between fabric and zone set at selected data source *Url entity*.

Explanation

Start enumeration collection.

Action

No action is required.

BTAFM0657I Start enumerating entity of association between fabric and zone at selected data source *Url entity*.

Explanation

Start enumeration collection.

Action

No action is required.

BTAFM0658I Start enumerating entity of association between fabric and zone alias at selected data source *Url entity*.

Explanation

Start enumeration collection.

Action

No action is required.

BTAFM0659I Start enumerating entity of association between switch and zone set at selected data source *Url entity*.

Explanation

Start enumeration collection.

Action

No action is required.

BTAFM0660I Start enumerating entity of association between switch and zone at selected data source *Url entity*.

Explanation

Start enumeration collection.

Action

No action is required.

BTAFM0661I Start enumerating entity of association between switch and zone alias at selected data source *Url entity*.

Explanation

Start enumeration collection.

Action

No action is required.

BTAFM0662I Start enumerating associations between virtual fabric and zoning entities at selected data source *Url entity*.

Explanation

Start enumeration collection.

Action

No action is required.

BTAFM0663I Starting collection of switch control processor blades for switch *switch identifier*.

Explanation

A probe job has begun to collect entities from the given switch.

Action

No action is required.

BTAFM0664I *Count of collected entities control processor blades collected for switch switch where entities exist.*

Explanation

A probe job has collected the given number of control processor blades so far.

Action

No action is required.

BTAFM0665I *Collection of switch control processor blades completed. Count of collected entities entities collected in total for switch switch identifier.*

Explanation

A probe job has collected the given number of entities.

Action

No action is required.

BTAFM0666I *Checksums for the active and defined Zone Database could not be updated for fabric entity.*

Explanation

Zone Database Checksums are not available for the fabric, or an error occurred while processing the Checksums. These Checksums are used only as a performance enhancement to quickly determine whether the Zone Database on the fabric has changed. This message is expected in certain scenarios. Since the Checksums cannot be saved, the Zone Database will be collected from the SMI-S provider during the next Probe.

Action

No action is required.

BTAFM0667E *Job id or request id is missing for a SRA job that is been processed.*

Explanation

Server is attempting to process a job for which response was received from SRA. Job id and/or request id is missing and so the response cannot be matched with a server job that is currently running.

Action

No action is required.

BTAFM0668E *Command and/or job timestamp is missing for job id with request id .*

Explanation

Command and/or job timestamp is missing for a SRA job and so the job cannot be processed successfully.

Action

No action is required.

BTAFM0669I job *id* with *request id* was is not found. Device server may have been restarted after job was created.

Explanation

Job was notfound and will not be processed. Device server may have been restarted after the job was created.

Action

No action is required.

BTAFM0670E could not retrieve output file for job *id* with *request id* .

Explanation

Unable to retrieve output file from agent for specified job. Job will be considered failed.

Action

No action is required.

BTAFM0671E Another probe of fabric *The Name+Nameformat of the fabric* is already in progress.

Explanation

Another probe for the same fabrric was already started and is in progress, so the new probe cannot be started.

Action

Start again the new probe only after the previous one is finished.

BTAFM0672E Device server is not registered with agent manager. Will not be able to invoke scanner on host .

Explanation

The device server is unable to register with the agent manager using the credentials provided through the agent manager registration panel.

Action

Verify the credentials specfied through the agent manager registration configuration panel. Also verify the connectivity between hosts running device server and agent manager if they are been hosted on different servers.

BTAFM0673E There are no agents that are currently available to probe fabric .

Explanation

There are no agents configured to probe specified fabric or the agents configured are not operational.

Action

Configure agents to probe fabric. If agents are already configured, check if they are operational.

BTAFM0674W No fabric found for event source that is associated with switch with IP address .

Explanation

No fabric was found for event resource generated by the switch.

Action

No action is required.

BTAFM0675E Unable to start parsing of SRA fabric probe data for SRA job id request id file name .

Explanation

An error was encountered while attempting to start parsing of fabric probe data returned by SRA for specified job id and request id.

Action

Check if there are any connectivity issues with the SRA. Also, check the device server message and trace log for more detail. If the problem continues, contact IBM support.

Related reference

- [Getting support](#)

BTAFM0676E Error parsing SRA fabric probe data for SRA job id request id file name .

Explanation

An error was encountered while parsing of fabric probe data returned by SRA for specified job id and request id.

Action

Check the device server message and trace log for more detail. If the problem continues, contact IBM support.

Related reference

- [Getting support](#)

BTAFM0677E Unable to connect to SNMP port (another application may already be connected and forwarding messages) .

Explanation

The device server is trying to connect to port for all SNMP Traps.

Action

Check to make sure another process is not listen on the port.

Related reference

- [Getting support](#)

BTAFM0678I The *Name of the switch* switch was removed.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BTAFM0679I The *The Name+Nameformat of the fabric* fabric was removed.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BTAFM0680E The *Name of the switch* switch was not removed because it is not missing.

Explanation

There are Data Sources configured managing this switch. Only switches which are no more managed by any Data Source can be removed.

Action

Remove all Data Sources which are managing this switch.

BTAFM0681E The *The Name+Nameformat of the fabric* fabric was not removed because it is not missing.

Explanation

There are Data Sources configured managing switches of this fabric. The fabric can only be removed after all Data Sources which are managing switches of this fabric are removed.

Action

Remove all Data Sources which are managing switches of this fabric.

BTAFM0682E An error occurred while checking for access to the database to save new zoning information for fabric to the database.

Explanation

The device server is unable to save new zoning information to the database because an error occurred while checking for other currently running jobs that might cause conflicts.

Action

Try the operation again later. If the problem continues, restart the device server.

BTAFM0683E Unable to access the database to save zoning information for fabric . Another job is currently saving new

zoning information to the database for the same fabric.

Explanation

The device server cannot save new zoning information to the database because another job is saving new zoning information to the database for the same fabric. A maximum wait timeout occurred while attempting to access the database to save new zoning information for this fabric.

Action

Try the operation again later. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

BTAFM0684I The job is waiting to access the database to save new zoning information for fabric . Another job is currently saving zoning information to the database for the same fabric.

Explanation

The device server is waiting to save new zoning information to the database because another job is currently saving new zoning information for the same fabric. The job will continue when access to the database is available.

Action

No action is required.

BTAFM0685W Host/IP Address is not a switch.

Explanation

Host or IP address is not a switch.

Action

Please make sure to type in the correct host or IP address.

BTAFM0686W Switch is not a supported switch.

Explanation

Switch being processed is not a supported switch.

Action

Please make sure to specify a supported switch.

BTAFM0687W The switch does not respond to SNMP queries.

Explanation

Switch does not respond to SNMP queries.

Action

Please make sure to specify a valid switch. Check to see if the switch is configured to respond to SNMP queries.

BTA FM0688W Cannot communicate with host or IP address .

Explanation

Not able to communicate with the specified host or IP address.

Action

Please make sure to specify a valid fully qualified hostname or IP address.

BTA FM0689W No ports were discovered for the switch .

Explanation

The probe did not discover any ports for the switch. The switch cannot be monitored for performance information.

Action

Make sure the switch is correctly configured with the SMI agent.

BTA FM0690I Collection of data from trunks is completed. Data was collected from *count of collected entities* trunks.

Explanation

A probe job has collected data from trunks in your storage environment.

Action

No action is required.

BTA FM0691I Starting collection of data from trunks for switch *switch identifier*.

Explanation

A probe job has begun to collect entities from the given switch.

Action

No action is required.

BTA FM0692I *Count of entities* trunks collected for switch *switch where entities exist*.

Explanation

A probe job has collected the given number of trunks so far.

Action

No action is required.

BTA FM0692E A response from the data collector was not received within the specified time.

Explanation

The data collector did not respond to the server in the allotted time. The data collector might not be running or it might not be able to connect to the server.

Action

Verify that the data collector is running and that it can connect to the server. Then try the request again.

BTAFM0693E A response from the data collector was not received.
The request failed with return code *return_code*

Explanation

The specified scan did not complete.

Action

None.

BTAFM0694W Zoning data cannot be collected because there is a transaction in progress on the switch *key*

Explanation

The specified scan did not complete.

Action

None.

BTAFM0695E The switch *key* is returning unexpected data.

Explanation

The specified scan did not complete.

Action

None.

BTAFM0696E Zone set *zoneset_name* is already active.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric because a zone set with the specified name is already active.

Action

Try the operation again with a different zone set name.

BTAFM0697E Zone set *zoneset_name* is already inactive.

Explanation

A failure occurred during an attempt to configure zoning on the specified fabric because a zone set with the specified name is already inactive.

Action

Try the operation again with a different zone set name.

BTAFM0698E On the switch *switch_name* VSAN *vsan_name* was not found.

Explanation

The VSAN specified as a parameter can not be found in the switch.

Action

Try the operation again with a different VSAN name.

BTAFM0699E The switch *key* did not return zoning data.

Explanation

The switch did not return zoning data. There may be issues with the switch configuration or zoning configuration.

Action

Review recent configuration changes and try to run the probe again.

BTAFM0700E Duplicate entries for the same switch: *switch*.

Explanation

The list of switches contains duplicate entries. Either a DNS name matches an IP address or there are duplicate DNS names or IP addresses.

Action

Ensure the list of switches contains only one entry per switch.

BTAFM0701E Current active full zone configuration is not synchronized with the zone configuration on the switch *switch_name* for VSAN *vsan_name* .

Explanation

The active zone configuration for the VSAN does not match the configuration on the switch. The zoning action will not be performed.

Action

Activate the zone set using the switch command line or management application. Activating the zone set on a switch synchronizes the active zone configuration for the VSAN with the zone configuration of the switch.

BTAFM0702E You cannot monitor Brocade Access Gateway switches without Network Advisor.

Explanation

You must use Network Advisor to monitor Brocade Access Gateway switches, also known as NPV switches. You cannot use SNMP agents as the data source for Brocade Access Gateway switches, you must use SMI agents that are installed with Network Advisor.

Action

Use Network Advisor to monitor Brocade Access Gateway switches.

BTA FM0703I Waiting for probes of other Access Gateway switches to complete.

Explanation

An Access Gateway switch is a Brocade switch that is operating in Access Gateway mode. Only one Access Gateway switch at a time can be probed by using the same SMI agent.

A probe of an Access Gateway switch was requested, but a probe of another Access Gateway switch is already in progress. The current probe waits until the first probe is completed. If there are no probes of other Access Gateway switches in progress, the current probe proceeds without any wait.

Brocade switches in Access Gateway mode are also known as NPV switches.

Action

No action is required.

BTA FM0704W Distributing zone configuration across all the switches for VSAN *vsan_name* did not succeed on the switch *switch_name* .

Explanation

The active zone configuration for the VSAN does not match the configuration on the switch. The zoning action will not be performed.

Action

Activate the zone set using the switch command line or management application. Activating the zone set on a switch synchronizes the active zone configuration for the VSAN with the zone configuration of the switch.

BTA FM0705W Zone data collection after zone changes were made failed on the switch *switch_name* .

Explanation

Zone changes were successful. Zone data collection to update the latest zone changes failed. The most current zone information will not be reported.

Action

Run a data collection for the failing switch. If the problem persists, contact IBM Software Support.

BTA FM0706E The fabric probe was unable to collect some details of the blades on the switches.

Explanation

A problem was encountered during the probe of the switch blades. The switch blade properties might be reported incorrectly until a later probe is completed.

Action

Try the following actions to resolve the problem:

- Check the logs for an indication of the error or exception, and resolve the problem if possible. For information about probe logs, see the Viewing probe logs topic in the IBM Spectrum Control Knowledge Center.
- Try the probe again at a later time. For instructions on how to start a probe, see the Starting probes topic in the IBM Spectrum Control Knowledge Center.
- If the problem continues, contact IBM customer support.

BTA FM0707I You cannot use IBM Spectrum Control to make zoning changes for provisioning on switch *switch_name*.

Explanation

A check of the ability to perform zone control operations on the switch has failed. Zone control operations through IBM Spectrum Control are not possible on this switch. This affects the Provisioning feature.

Action

You can ignore this message if:

- You do not plan to use the provisioning feature of IBM Spectrum Control.
- You plan to use the provisioning feature of IBM Spectrum Control for volume creation and assignment but not for zoning.

Otherwise, check that the model of the switch supports fibre channel zone control, and check that the switch is configured properly for zone control.

BTAFM0708E The probe was unable to collect some details of the switches.

Explanation

A problem was encountered during the probe. The fabric or switch properties might be reported incorrectly until a later probe is completed.

Action

Try the following actions to resolve the problem:

- Check the logs for an indication of the error or exception, and resolve the problem if possible. For information about probe logs, see the Viewing probe logs topic in the IBM Spectrum Control Knowledge Center.
- Try the probe again at a later time. For instructions on how to start a probe, see the Starting probes topic in the IBM Spectrum Control Knowledge Center.
- If the problem continues, contact IBM customer support.

BTAFM0709I Started to process information for fabric *fabric_name*.

Explanation

The information is being parsed and added to the database.

Action

No action is required.

BTAFM0710I Started to process information for switch *switch_name*.

Explanation

The information is being parsed and added to the database.

Action

No action is required.

BTAFM0711I Started to process information for discovered switches.

Explanation

The information is being parsed and added to the database.

Action

No action is required.

BTAFM0712I Started to process information for a switch blade.

Explanation

The information is being parsed and added to the database.

Action

No action is required.

BTAFM0713I Started to process information for a switch zone set.

Explanation

The information is being parsed and added to the database.

Action

No action is required.

BTAFM0714I Started to process information for switch ports.

Explanation

The information is being parsed and added to the database.

Action

No action is required.

BTAFM0715E Error occurred while processing information for fabric *fabric_name*.

Explanation

The information for the fabric was not added to the database.

Action

Check the Device server message log and trace log for more details. If the problem continues, contact IBM Support.

BTAFM0716E Error occurred while processing information for virtual fabric *virtual_fabric_name*.

Explanation

The information for the virtual fabric was not added to the database.

Action

Check the Device server message log and trace log for more details. If the problem continues, contact IBM Support.

BTAFM0717E Error occurred while processing information for switch *switch_name*.

Explanation

The information for the fabric and switch was not added to the database.

Action

Check the Device server message log and trace log for more details. If the problem continues, contact IBM Support.

BTAFM0718E Error occurred while processing information for discovered switches.

Explanation

The information for the discovered switches was not added to the database.

Action

Check the Device server message log and trace log for more details. If the problem continues, contact IBM Support.

BTAFM0719E Error occurred while processing information for logical switches.

Explanation

The information for the logical switch was not added to the database.

Action

Check the Device server message log and trace log for more details. If the problem continues, contact IBM Support.

BTAFM0720E Error occurred while processing information for active zone set *active_zone_set_name*.

Explanation

The information for the active zone set was not added to the database.

Action

Check the Device server message log and trace log for more details. If the problem continues, contact IBM Support.

BTAFM0721E Error occurred while processing information for inactive zone set *inactive_zone_set_name*.

Explanation

The information for the inactive zone set was not added to the database.

Action

Check the Device server message log and trace log for more details. If the problem continues, contact IBM Support.

BTAFM0722E Error occurred while processing information for port *port_name*.

Explanation

The information for the port was not added to the database.

Action

Check the Device server message log and trace log for more details. If the problem continues, contact IBM Support.

BTAQE1107E InbandScanHandler failed to start InbandScanner *scanner name* on managed host *target*.

Explanation

The InbandScanner scans managed hosts for device information and displays that information in the network topology display. The InbandScanner is necessary for providing accurate information about your SAN.

Action

Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)

BTAQE1108E InbandScanHandler failed to get callback information for InbandScanner *scanner name* on managed host *target*.

Explanation

The InbandScanHandler service was unable to receive device information from the scanner. This can cause the network topology displays to show inaccurate SAN data. Other services will be affected as well.

Action

Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)

BTAQE1112E During an outband scan, the scanner *scanner name* was unable to identify the target host *target*.

Explanation

The target host might have an invalid IP address.

Action

Verify the IP address of the target managed host. Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)

BTAQE1113E Unable to invoke an Outband scan *scanner name* on target *target*.

Explanation

IBM Spectrum Control was unable to start an outband scan for Fabric. There might be problems with the SAN connectivity.

Action

Make sure the SAN is working properly. Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)

BTAQE114E OutbandScannerHandler received invalid callback information for Outband scanner *scanner name* on target *target*.

Explanation

The information received from an outband scan appears to be invalid.

Action

Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)

BTAQE1115E The outband scanner *scanner name* did not return the SAN ID on target *target*.

Explanation

The switch vendor may not support FE MIB or certain fields in the FC MGMT MIB.

Action

Check the Tivoli Support Web site to see if the switch is supported by IBM Spectrum Control. Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)

BTAVM0001I The operation *Name of the operation* processed successfully.

Explanation

The operation on the Virtualization Manager completed successfully. No error condition encountered.

Action

No action required.

BTAVM0002I The Web service call *Name of the operation* processed successfully.

Explanation

The Web service call from the Virtualization Manager completed successfully. No error condition encountered.

Action

No action required.

BTAVM0003I Data source *Name of the datasource* successfully added.

Explanation

The data source has been added successfully.

Action

No action required.

BTAVM0004I Data source *Name of the datasource* successfully deleted.

Explanation

The data source has been deleted successfully.

Action

No action required.

BTAVM0005I Data source *Name of the datasource* successfully modified.

Explanation

The data source has been modified successfully.

Action

No action required.

BTAVM0006I Discovery on data source *Name of the datasource* has started.

Explanation

The discovery on the data source collects the hypervisors managed by that data source.

Action

No action required.

BTAVM0007I Discovery on data source *Name of the datasource* completed successfully.

Explanation

The discovery on the data source completed successfully.

Action

No action required.

BTAVM0008I Probe of hypervisor *Name of the Hypervisor* has started.

Explanation

The probe collects configuration details of the hypervisor.

Action

No action required.

BTAVM0009I Probe of hypervisor *Name of the Hypervisor* completed successfully.

Explanation

The probe of the hypervisor completed successfully.

Action

No action required.

BTAVM0010I A connection test to data source *Name of the data source* has started.

Explanation

A connection test to the data source has started.

Action

No action required.

BTAVM0011I The Connection test to data source *Name of the data source* completed successfully.

Explanation

Connectivity to the data source was validated successfully.

Action

No action required.

BTAVM0012I Hypervisor *Name of the Hypervisor* discovered/rediscovered.

Explanation

The hypervisor has been discovered.

Action

No action required.

BTAVM0013I Discovery: Hypervisor *Name of the hypervisor* will not be discovered as it is managed by another data source.

Explanation

The hypervisor is managed by another data source.

Action

No action is required.

BTAVM0014I Discovery: Hypervisor *Name of the hypervisor* will not be discovered as it itself is registered as a data source.

Explanation

The hypervisor is registered directly as a data source.

Action

No action is required.

BTAVM0015I Collection of the physical storage configuration for hypervisor *Name of the hypervisor* has started.

Explanation

The probe collects physical storage configuration details of the hypervisor.

Action

No action is required.

BTAVM0016I Collection of the physical storage configuration for hypervisor *Name of the hypervisor* completed successfully.

Explanation

The probe of the hypervisors physical storage configuration completed successfully.

Action

No action is required.

BTAVM0017I Collection of the logical storage configuration for hypervisor *Name of the hypervisor* has started.

Explanation

The probe collects logical storage configuration details of the hypervisor .

Action

No action is required.

BTAVM0018I Collection of the logical storage configuration for hypervisor *Name of the hypervisor* completed successfully.

Explanation

The probe of the hypervisors logical storage configuration completed successfully.

Action

No action is required.

BTAVM0019I Collection of the virtual machines configuration for hypervisor *Name of the hypervisor* has started.

Explanation

The probe collects virtual machines configuration details of the hypervisor.

Action

No action is required.

BTAVM0020I Collection of the virtual machines configuration for hypervisor *Name of the hypervisor* completed successfully.

Explanation

The probe of the hypervisors virtual machines configuration completed successfully.

Action

No action is required.

BTAVM0021I The probe of *name of the hypervisor* found *number of physical disks* physical disks.

Explanation

The probe collects the number of physical disks found in the hypervisors physical storage configuration.

Action

No action is required.

BTAVM0022I The probe of *name of the hypervisor* found *number of logical volumes* logical volumes.

Explanation

The probe collects the number of logical volumes found in the hypervisors logical storage configuration.

Action

No action is required.

BTAVM0023I The probe of *name of the hypervisor* found *number of virtual machines* virtual machines.

Explanation

The probe collects the number of virtual machines found in the hypervisors virtual machines configuration.

Action

No action is required.

BTAVM0024I The *Name of the hypervisor* hypervisor was removed.

Explanation

This message is for informational purposes only.

Action

No further action is required.

BTAVM0025I VMWare Cluster *Name of the Cluster* discovered/rediscovered.

Explanation

The cluster has been discovered.

Action

No action required.

BTAVM1301I The probe of *name of the hypervisor* could collect partial information only for the disk with the device name *Device name of the disk*.

Explanation

The disk attributes number of heads, number of sectors and number of cylinders are not available.

Action

No action is required.

BTAVM1302I LUN correlation is not supported for disk with device name *Device name of the disk*, vendor: *Vendor name*, model: *model name*, for hypervisor *hypervisor name*.

Explanation

LUN correlation is not supported for this disk. LUN definition data for this disk will not be available.

Action

Refer to the IBM Spectrum Control Supported Product List.

BTAVM1503E An internal error occurred: *Text describing the internal error.*

Explanation

An internal operating error occurred. Any exceptions are logged in the traceTPCDeviceServer.log file. This log file resides in the installation subdirectory device/log.

Action

Review the traceTPCDeviceServer.log file. If the problem persists, contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2001E The mandatory parameter *Name of the mandatory parameter which is missing* is missing.

Explanation

A mandatory parameter for a Virtualization Manager operation is missing.

Action

Contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2002E Invalid parameter *Name of the parameter which was invalid.*

Explanation

A parameter passed to the Virtualization Manager is invalid.

Action

Contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2003E A database error was encountered during database query or insert.

Explanation

Data could not be retrieved from or inserted into the database. There may be a problem with database access or the database is not available.

Action

Try the action again. If the problem persists, check the Device server log files for error messages that might help determine the problem.

BTAVM2004E Cannot connect to the database repository.

Explanation

IBM Spectrum Control cannot connect to the database repository. There may be a problem with database access or the database is not available.

Action

Verify that the database service is up and running. Verify that you have a network connection to the computer on which the database repository is located. Try the action again.

BTAVM2006E The operation *Name of the operation that failed* failed for the following reason: *Reason of the failure*.

Explanation

The operation on the Virtualization Manager failed.

Action

Contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2007E The Web service call *Name of the operation* failed for the following reason: *Reason of the failure*.

Explanation

The Web service call from the Virtualization Manager has failed.

Action

Contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2008E The product *Name of the unsupported product* is not supported.

Explanation

This virtualization product is not supported by IBM Spectrum Control.

Action

Use a supported virtualization product. Refer to the IBM Spectrum Control Supported Product List.

BTAVM2010E The user name or password is invalid for *Address of the host*

Explanation

The user name or password is not valid.

Action

Enter the correct user name and password. Try the action again.

BTAVM2011E The operation *Name of the timed out operation* could not complete within the time limit of *Timeout threshold in milliseconds* milliseconds.

Explanation

For each service method invocation a timeout is defined. The timeout for this operation has been exceeded.

Action

Check the Hypervisor or increase the timeout value.

BTAVM2012E An error occurred while trying to establish secure communication over SSL.

Explanation

The communication to the data source over SSL failed. There may be a configuration error.

Action

Make sure that the SSL certificate of the data source has been imported correctly to the local truststore. If the problem persists, contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2013E The Add Device wizard could not add the *Name of the data source* data source.

Explanation

The VMWare data source might already be monitored by IBM Spectrum Control. A data source that is being monitored cannot be added again.

Action

Ensure that the data source is not already being monitored by IBM Spectrum Control. Verify that the IBM Spectrum Control server is up and running and that you have a network connection. If the data source is not being monitored, try to add it again. Ensure that you enter the correct authentication credentials in the wizard.

BTAVM2014E The deletion of data source *Name of the data source* failed.

Explanation

Failed to delete the data source from the database.

Action

Contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2015E The modification of data source *Name of the data source* failed.

Explanation

Failed to modify the data source in the database. The data source was not modified.

Action

Contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2016E Discovery on data source *Name of the datasource* failed.

Explanation

The discovery on the data source failed.

Action

Check the trace logs to find the error reason. The log files reside in the installation subdirectory device/log. If possible correct the error and rerun the discovery on the data source. If the problem persists, contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2017E Probe of the hypervisor *Name of the Hypervisor* failed.

Explanation

The probe of the hypervisor failed.

Action

Check the trace logs to find the error reason. The log files reside in the installation subdirectory device/log. If possible correct the error and rerun the probe of the hypervisor. If the problem persists, contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2018E IBM Spectrum Control can't connect to the data source *Name of the datasource*.

Explanation

A connection test to the data source failed.

Action

Check the trace logs to find the error reason. The log files reside in the installation subdirectory device/log. If possible correct the error and rerun the connection test to the data source. If the problem persists, contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2201E Probe: An error occurred during the collection of the physical storage configuration.

Explanation

The physical storage configuration of the hypervisor could not be collected.

Action

Check the configuration of your environment. Check the trace logs. The log files reside in the installation subdirectory device/log. If possible correct any error and rerun the probe of the hypervisor. If the problem persists, contact IBM Technical Support.

Related reference

-  [Getting support](#)

BTAVM2202E Probe: An error occurred during the collection of the logical storage configuration.

Explanation

The logical storage configuration of the hypervisor could not be collected.

Action

Check the configuration of your environment. Check the trace logs. The log files reside in the installation subdirectory device/log. If possible correct any error and rerun the probe of the hypervisor. If the problem persists, contact IBM Technical Support.

Related reference

-  [Getting support](#)

BTAVM2204E Probe: An error occurred during the collection of the virtual machine configuration.

Explanation

The virtual machines on the hypervisor could not be collected.

Action

Check the configuration of your environment. Check the trace logs. The log files reside in the installation subdirectory device/log. If possible correct any error and rerun the probe of the hypervisor. If the problem persists, contact IBM Technical Support.

Related reference

-  [Getting support](#)

BTAVM2206E Discovery: the hypervisor *Name of the hypervisor* will not be discovered because its version is not supported.

Explanation

The IBM Spectrum Control does not support the version. Refer to the IBM Spectrum Control Supported Product List.

Action

Refer to the IBM Spectrum Control Supported Product List and upgrade the Hypervisor code level to a supported level.

BTAVM2207E Calculation of the summary data for the hypervisor *Name of the hypervisor* failed.

Explanation

The calculation of the summary information for the hypervisor did not succeed.

Action

Check the configuration of your environment. Check the trace logs. The log files reside in the installation subdirectory device/log. If possible correct any error and rerun the probe of the hypervisor. If the problem persists, contact IBM Technical Support.

Related reference

-  [Getting support](#)

BTAVM2208E Unable to obtain the hypervisor version(s) from the datasource *Name of the datasource*.

Explanation

The discovery was not able to obtain the version of the hypervisors.

Action

Check the configuration of your environment. Check the trace logs. The log files reside in the installation subdirectory device/log. If possible correct any error and rerun the probe of the hypervisor. If the problem persists, contact IBM Technical Support.

Related reference

-  [Getting support](#)

BTAVM2209E Unable to obtain information about other Virtual Centers managing the hypervisor(s) of datasource *Name of the datasource*.

Explanation

The discovery was not able to obtain information about other Virtual Centers which are managing the hypervisor(s).

Action

Check the configuration of your environment. Check the trace logs. The log files reside in the installation subdirectory device/log. If possible correct any error and rerun the probe of the hypervisor. If the problem persists, contact IBM Technical Support.

Related reference

-  [Getting support](#)

BTAVM2210W Error getting LUN definition data for the disk with the device name *Device name of the disk*, storage subsystem vendor: *Vendor name*, model: *model name*, for hypervisor *hypervisor name*.

Explanation

The LUN definition data will not be available.

Action

Check the configuration of your environment. Check the trace logs. The log files reside in the installation subdirectory device/log. If possible correct any error and rerun the probe of the hypervisor. If the problem persists, contact IBM Technical Support.

BTAVM2211E Probe: Virtualization Manager failed to get the VMWare VI data source for the hypervisor *Name of the hypervisor* from the database.

Explanation

The VMWare VI data source managing the hypervisor may have been deleted.

Action

Add the VMWare VI data source, perform a VMWare VI data source discovery and re-create the probe.

BTAVM2212E Probe: The hypervisor *Name of the hypervisor* is not available on the VMWare VI datasource *Name of the datasource*.

Explanation

The hypervisor may have been deleted or moved to another datasource.

Action

Add the VMWare VI datasource the hypervisor was moved to. Perform a VMWare VI data source discovery and re-create the probe.

BTAVM2213E Data source *Name of the datasource* is disconnected from Virtual Center.

Explanation

A connection test to the data source failed.

Action

Connect data source in Virtual Center and run discovery again.

BTAVM2214E The probe job encountered an NFS file system while probing ESX server {0}. IBM Spectrum Control currently does not support probes of ESX servers with NFS file systems. The probe job for this ESX server has been stopped. Probes of other ESX servers that are included in this probe job will continue.

Explanation

ESX servers with NFS file systems are not supported. Probes of other ESX servers that do not have NFS file systems will continue.

Action

If this is a repeating probe job, remove the ESX servers that have NFS file systems from the probe job to prevent this error from occurring again.

BTAVM2215W Unsupported storage subsystem disk with device name *Device name of the disk*, vendor: *Vendor name*, model: *model name*, for hypervisor *hypervisor name* with hypervisor version less than 3.5.0.

Explanation

Hypervisors with version less than 3.5.0 do not offer the necessary support for getting LUN definition data for this storage subsystem type. LUN definition data for this disk will not be available.

Action

Refer to the IBM Spectrum Control Supported Product List and upgrade the Hypervisors and Virtual Centers to the latest supported level.

BTAVM2216E Unable to get keystore instance.

Explanation

The required keystore type (JKS) is not available.

Action

Make sure the required keystore application is available in the Java environment. If the problem persists, contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2217E Unable to load keystore file.

Explanation

The default keystore file used by IBM Spectrum Control could not be loaded.

Action

Make sure that the keystore password has not been changed. If the problem persists, contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2218E Unable to set certificate entry in keystore file.

Explanation

The certificate entry couldn't be added to the keystore file

Action

Make sure the certificate being added to the keystore is a valid certificate. If the problem persists, contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2219E Unable to open keystore for writing.

Explanation

An exception occurred attempting to open keystore for writing.

Action

Make sure that the keystore file is valid and has the proper permissions. If the problem persists, contact IBM Technical Support.

BTAVM2220E Unable to close keystore file.

Explanation

An exception occurred attempting to close keystore file.

Action

Make sure there are proper permissions on the keystore file and that no other user or process has it locked.

BTAVM2221E Unable to acquire lock on keystore file.

Explanation

The keystore file cannot be locked for write access.

Action

Make sure that no other user or process has the keystore locked.

BTAVM2222E Unable to store certificate in keystore file.

Explanation

The certificate could not be stored in the keystore file.

Action

Make sure that the keystore is valid and that the keystore password has not been changed.

BTAVM2223E Unable to release lock on keystore file.

Explanation

Keystore could not be unlocked.

Action

Make sure that the keystore is accessible and no other user or process has it locked.

BTAVM2224E Unable to decrypt keystore password.

Explanation

An error occurred while trying to decrypt the provided password to the keystore.

Action

The keystore may have become corrupted. Save the keystore to a new location so a new one will be created in the default location, and try again. If the problem persists, contact IBM Technical Support.

Related reference

-  [Getting support](#)

BTAVM2225E Unable to open keystore for reading.

Explanation

The keystore could not be opened for reading.

Action

Make sure that the keystore is not locked and is set with the proper read permissions.

BTAVM2226E Certificate already exists in keystore.

Explanation

Unable to store the certificate in the keystore because a certificate already exists for hostname.

Action

Either choose the option to replace the certificate, or specify a different alias (hostname) for the certificate.

BTAVM2227E *host_address* hypervisor is already being monitored and could not be added.

Explanation

A hypervisor cannot be added if it is already monitored by IBM Spectrum Control.

Action

Verify that the hypervisor has already been added to IBM Spectrum Control. If the hypervisor is not being monitored, try to add it again. If it is being monitored, enter the name of a different hypervisor.

BTAVM2228E Missing host name.

Explanation

The host name of the hypervisor was not provided.

Action

Make sure to provide a hostname for the certificate.

BTAVM2229E Missing certificate.

Explanation

A valid certificate was not provided

Action

Make sure to provide a valid certificate for insertion into the keystore.

BTAVM2230E Cannot create keystore directory.

Explanation

The default keystore directory could not be created.

Action

Make sure file permissions have been set to allow creation of the default keystore directory.

BTAVM2231E Cannot download the certificate from Data Source *Name of the data source*.

Explanation

The SSL certificate cannot be downloaded from the specified data source. The certificate must be downloaded before the data source can be added. A data source can be a VMWare hypervisor or vSphere server.

Action

Verify that the specified Data Source is a VMWare hypervisor or vSphere server. Verify that you have a network connection to the data source. Try to add the data source again.

BTAVM2232E Cannot connect to the *Name of the data source* data source.

Explanation

The host name or IP address that was entered for the data source might not be valid.

Action

Ensure that the network is up and available and that you have a network connection to the data source. Verify the correct host name or IP address. Try to add the data source again.

BTAVM2233E Cannot download the certificate from the port.

Explanation

The port that was entered for the data source might not be valid.

Action

Specify the correct port. The default port number is 443.

BTAVM2234E The *hypervisor name* hypervisor was not removed because IBM Spectrum Control is running other actions on the device.

Explanation

The hypervisor is in use by IBM Spectrum Control and cannot be removed from the database repository at this time. For example, a probe schedule is collecting data about the hypervisor.

Action

Wait for the probe schedule or other action to complete and try to remove the hypervisor again.

BTAVM2235E Unable to obtain the cluster(s) from the datasource *Name of the datasource*.

Explanation

The discovery was not able to obtain the clusters.

Action

Check the configuration of your environment. Check the trace logs. The log files reside in the installation subdirectory device/log. If possible correct any error and rerun the discovery of the hypervisor(s). If the problem persists, contact IBM Technical Support.

Related reference

- [Getting support](#)

BTAVM2236W Subsequent steps of probe process may not be able to collect data for the hypervisor *Name of the hypervisor* because the hypervisor is in critical state.

Explanation

The hypervisor is in critical state and the probe process may not be able to do data collection.

Action

Check the VMWare ESX or VirtualCenter for the cause of critical state for the hypervisor. If possible correct the issue and rerun the probe of the hypervisor.

BTAVM2237E Datastore Browser Task failed for hypervisor *Name of the hypervisor*, datastore *Name of the datastore* with error: *Error*

Explanation

The datastore could not be browsed.

Action

The datastore could not be browsed due to the error specified. If possible correct the issue and rerun the probe of the hypervisor.

BTAVM2238E The registration of the vSphere Web Client extension for IBM Spectrum Control has started on *Name of the vCenter server*.

Explanation

The registration process has started.

Action

No action required.

BTAVM2239E The registration of the vSphere Web Client extension for IBM Spectrum Control did not extract the extension package.

Explanation

The configuration process updates the vSphere Web Client extension for IBM Spectrum Control Centre package. The extension package, TPC_VmPlug.zip, could not be updated because the packaged was not extracted.

Action

Check the trace logs. The log files reside in the installation subdirectory log under device. Correct any error and rerun the discovery of the hypervisor(s). If the problem persists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTAVM2240E The registration of the vSphere Web Client extension for IBM Spectrum Control did not complete while updating the VASA web archive file, vasa.war, with the IBM Spectrum Control server configuration.

Explanation

The configuration process updates the IBM Spectrum Control Web Client extension for VMware package. The update operation could not be completed.

Action

Check the trace logs. The log files reside in the installation subdirectory log under device. Correct any error and rerun the discovery of the hypervisor(s). If the problem persists, contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTAVM2241E The registration of the vSphere Web Client extension for IBM Spectrum Control completed.

Explanation

The registration process has completed.

Action

No action required.

BTAVM2242E Unable to register IBM Spectrum Control as an extension on the vCenter server *Name of the vCenter server*. The validation of input values did not complete.

Explanation

The validation of input values did not complete. One or more of the supplied values may be empty or invalid.

Action

Check the trace logs. The log files reside in the installation subdirectory log under device. Correct any error and rerun the discovery of the hypervisor(s). If the problem persists, contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTAVM2243E Unable to register IBM Spectrum Control as an extension on the vCenter server *Name of the vCenter server*. Could not authenticate with the vCenter server.

Explanation

Could not authenticate with the vCenter server.

Action

Check the trace logs. The log files reside in the installation subdirectory log under device. Correct any error and rerun the discovery of the hypervisor(s). If the problem persists, contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

BTAVM2244E The registration of the vSphere Web Client extension for IBM Spectrum Control did not complete.

Explanation

The configuration process was unable to register the IBM Spectrum Control Web Client extension for VMware

Action

Check the trace logs. The log files reside in the installation subdirectory log under device. Correct any error and rerun the discovery of the hypervisor(s). If the problem persists, contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

BTAVM2245E Unable to connect to the vCenter Server *Name of the datasource*.

Explanation

The data source configuration process was not able to connect to the vCenter server.

Action

Check the trace logs. The log files reside in the installation subdirectory log under device. Correct any error and rerun the discovery of the hypervisor(s). If the problem persists, contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

BTAVM2246E Unable to configure the vCenter Server.

Explanation

The configuration process was unable to remotely register the vSphere Web Client extension for IBM Spectrum Control.

Action

Check the trace logs. The log files reside in the installation subdirectory log under device. Correct any error and rerun the discovery of the hypervisor(s). If the problem persists, contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

BTAVM2247E The registration of the vSphere Web Client extension for IBM Spectrum Control did not delete the temporary directory *Name of the directory*.

Explanation

The registration process was unable to delete a temporary directory on the vCenter Server

Action

Check the trace logs. The log files reside in the installation subdirectory log under device. Correct any error and rerun the discovery of the hypervisor(s). If the problem persists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTAVM2248E The registration of IBM Spectrum Control as a VASA provider did not complete.

Explanation

The configuration process was unable to register IBM Spectrum Control as a VASA provider for VMware.

Action

Check the trace logs. The log files reside in the installation subdirectory log under device. Correct any error and rerun the discovery of the hypervisor(s). If the problem persists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTAVM2249E Automatic registration of IBM Spectrum Control as a VASA provider is not supported for vCenter Server version 5.0 and earlier.

Explanation

vCenter Server version 5.1 or later is required to automatically register IBM Spectrum Control as a VASA provider.

Action

Register IBM Spectrum Control as a VASA provider manually in the vSphere Web Client.

BTAVM2250E IBM Spectrum Control is already registered as a VASA provider for vCenter Server *server_name*. Register IBM Spectrum Control as a VASA provider manually in the vSphere Web Client to update the credentials.

Explanation

Automatic update of IBM Spectrum Control as a VASA provider is not supported if the provider is already registered for this vCenter Server.

Action

Remove IBM Spectrum Control as a VASA provider and register it again manually to update the user name or password.

BTAVM2251E One or more third-party VASA providers are already registered with the vCenter Server. IBM Spectrum Control VASA provider was not registered. Register IBM Spectrum Control as a VASA provider manually.

Explanation

If one or more third-party VASA providers are already registered with the vCenter Server, you must register IBM Spectrum Control VASA provider manually.

Action

Register IBM Spectrum Control as a VASA provider manually in the vSphere Web Client.

BTAVM2252E The registration of IBM Spectrum Control as a VASA provider has started on *Name of the vCenter server*.

Explanation

The VASA provider registration process has started.

Action

No action required.

BTAVM2253E The registration of IBM Spectrum Control as a VASA provider has completed.

Explanation

The VASA provider registration process has completed.

Action

No action required.

BTAVM2254E The registration of the vSphere Web Client extension for IBM Spectrum Control did not complete. The current session is invalid.

Explanation

The current session is invalid.

Action

Log out of the vCenter Server and then log back in.

BTAVM2255E The registration of IBM Spectrum Control as a VASA provider did not complete. The current session is invalid.

Explanation

The current session is invalid.

Action

Log out of the vCenter Server and then log back in.

BTAVM2256W Could not determine the host for VM with ID: *host id* and Name: *Vendor name*. Check if the same mac address is used on other computers.

Explanation

The correlation between the VM and the host cannot be done if the MAC Addresses are not unique.

Action

Change the MAC Address of the VM.

BTAVM2257I Found *number of files* files on *name of datastore* of *name of the hypervisor*.

Explanation

The probe collected details on given number of files from the datastore of the hypervisor.

Action

No action is required.

BTAVM2258I The probe of *name of the hypervisor* found *number of controllers* controllers.

Explanation

The probe collected the number of controllers for the hypervisors.

Action

No action is required.

BTAVM2259I Collecting file system details for hypervisor *Name of the hypervisor*.

Explanation

The probe collects file system details of the hypervisor.

Action

No action is required.

BTAVM2260I Collecting list of files for hypervisor *Name of the hypervisor*.

Explanation

The probe collects list of files on the datastores Yes of the hypervisor.

Action

No action is required.

BTAVM2261I Collecting logical volumes for hypervisor *Name of the hypervisor*.

Explanation

The probe collects logical volumes details of the hypervisor.

Action

No action is required.

BTAVM2262I Collecting disk partition for hypervisor *Name of the hypervisor*.

Explanation

The probe collects disk partitions details of the hypervisor.

Action

No action is required.

BTAVM2263I Files details for *Name of the datastore* being collected by *id of the Hypervisor*.

Explanation

The file details for the given datastore is already in progress as part of another hypervisor's probe, hence will be skipped.

Action

No action is required.

BTAVM2264I Files details for *Name of the datastore* were collected by *id of the Hypervisor* on *timestamp*.

Explanation

The file details for the given datastore were collected on given time as part of another hypervisor's probe, hence will be skipped.

Action

No action is required.

BTAVM2265E Invalid host name or IP address.

Explanation

Host name or IP address is not valid.

Action

Enter a valid host name or IP address

BTAVM2266E The connection information cannot be updated because it points to another device.

Explanation

To change the host name or IP address at least one hypervisor has to be managed by the new host name/IP address.

Action

Enter the correct host name or IP address.

BTAVM2268E The connection information cannot be updated because IBM Spectrum Control cannot determine if the hypervisor is managed by the *Name of the data source* data source.

Explanation

This problem might occur if the hypervisor is unavailable or the connection was lost.

Action

Try the following actions:

- Verify that the hypervisor is available.
- Ensure that a connection to the hypervisor is active.
- Check the log files for error messages to determine the cause. For information about the location of log files, see the IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93/>.

BTAVM2269E The connection information cannot be updated because a data source with this host name or IP address is already present.

Explanation

You cannot change the host name or IP address of a data source to match an existing one because you cannot move hypervisors from one data source to another data source.

Action

Let the existing event mechanism detect when a hypervisor was moved to another data source.

BTAVM2270E The connection information cannot be updated because it doesn't point to a data source of the same type (vCenter/ESX).

Explanation

Changing the host name or IP address of a data source is allowed only within the same connection type. You cannot change a vCenter Server address to an ESX Server address. Similarly, you cannot change an ESX Server address to a vCenter address.

Action

Enter a host name or IP address that points to a data source of the same type as the one you are trying to update.

BTAVM2271W The hypervisor *Name of the Hypervisor* cannot be discovered because its connection state is "*Connection State*".

Explanation

The vCenter Server discovery process ignores the hypervisors that do not have the connection state of "CONNECTED". The hypervisor might have been disconnected from the vCenter Server intentionally for maintenance purposes. It might have been moved to another vCenter Server, but remains disconnected in the inventory of the old vCenter Server. Or it might have been accidentally disconnected due to a functional problem.

Action

Verify why the hypervisor is not connected to the vCenter Server.

BTAVM2272E The user *User Name* does not have the privilege to browse the datastore *Name of the Datastore*.

Explanation

The hypervisor scan failed because the user does not have the Datastore.Browse privilege that is required to execute the SearchDatastoreSubFolders task. For more details, go to the IBM Knowledge Center and check Administering>Administering resources and data sources>Hypervisors and VMware data sources>Checking permissions to browse data stores. You can access the IBM Knowledge Center for the product at <http://www.ibm.com/support/knowledgecenter/SS5R93>.

Action

Assign the user to a role that has the Datastore.Browse privilege. An appropriate default role is Virtual machine power user.

HWN020001I *Operation Name of the operation processed successfully.*

Explanation

The request has been executed successfully. No error condition has been encountered.

Action

No action is required.

HWN020002E *Mandatory parameter Name of the mandatory parameter which is missing* missing

Explanation

Mandatory parameter {0} missing

Action

Pass the mandatory parameters.

HWN020003E *Invalid parameter Name of the parameter which was* *invalid*

Explanation

Invalid Parameter{0}

Action

Pass a valid parameter.

HWN020101E *The external process terminated unexpected.*

Explanation

The external process for the current job did not complete successful because the process terminated unexpected.

Action

If available check the logfile of the external process to find the cause of the error. Make sure that no external event on the IBM Spectrum Control server terminated or ended the process. Try to rerun the job. If the problem persists, enable high level tracing as explained in the IBM Spectrum Control Installation and Configuration Guide and contact IBM Software Support.

HWN020102W *The external process was canceled per users request.*

Explanation

The external process for the current job did not complete successful because the process was terminated per users request to cancel the job.

Action

No action is required.

HWN020103E The external process exceeded the timeout limit and was canceled.

Explanation

The external process for the current job did not complete within the timeout limit and was terminated.

Action

Check your system environment, including the network connection and server load to determine the root cause of the long running job. If necessary, increase the timeout value for the external process. For information about increasing the timeout value, go to the IBM Knowledge Center at http://www.ibm.com/support/knowledgecenter/search/fqz0_t_setting_timeout_values_for_the_deviceserver?scope=SS5R93. View the appropriate topic for the version of IBM Spectrum Control that you are using.

HWN020104E The external process could not be started.

Explanation

Unable to start the external process for the current job.

Action

If available check the logfile of the external process to find the cause of the error. Try to rerun the job. If the problem persists, enable high level tracing as explained in the IBM Spectrum Control Installation and Configuration Guide and contact IBM Software Support.

HWN020105E The data collector is not responding to the server.

Explanation

The data collector did not respond to the server in the allotted time. The data collector might not be running or it might not be able to connect to the server.

Action

Verify that the data collector is running and that it can connect to the server.

HWN020106E An external process was cancelled by the data collector.

Explanation

An external process for the current job did not complete because the process was cancelled during the shutdown of the data collector. The data collector was shut down for one of the following reasons:

- The user stopped the data collector.
- The system that the data collector runs on was shut down or restarted.
- The data collector is being upgraded.

Action

Wait until the data collector starts again, then try the task or job again.

HWN021503E The action cannot be completed

Explanation

The error occurred while processing a request from the GUI.

Action

To resolve the issue, try the following actions:

- Verify that the local area network is available and a firewall is not preventing network access to product services and agents.
- Check the status of the product servers and database repository on the Home > System Management page.
- On Windows, verify that the related database services are active.
- Check for error messages in the log files for the servers.

If the problem persists, check the product's log files for error messages that might help determine the problem. For information about the location of log files, check the IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/SS5R93/>.

Related reference

- [🔗 Troubleshooting problems with the IBM Spectrum Control component and servers](#)
- [🔗 Default locations of log files](#)

HWN021504E Entity *The ID of the entity* was not found.

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

The server or the interface might be out of sync with the co-server. Rerun the discovery process, and then run a data collection task.

HWN021508E Credentials not found

Explanation

There was a problem accessing the user credential on the coserver

Action

Ensure the database is running and that the IBM Spectrum Control setup was completed successfully.

HWN021514E The invocation of CIM method *Name of method* failed on SMI-S provider *Name of SMI-S provider* . The return code is *Return code of method*

Explanation

The extrinsic CIM method that was invoked on the given SMI-S provider failed.

Action

Ensure that the correct input parameters for the CIM method have been used.

HWN021515E The invocation of CIM method *Name of method* failed on SMI-S provider *Name of SMI-S provider* with the following exception text: *Exception text*

Explanation

The extrinsic CIM method that was invoked on the given SMI-S provider failed.

Action

Check the ErrorTrace.log file for further information. If the log file does not exist, enable tracing as follows: (1) Enter the WAS administration console. Click Application Servers, click the server name, and click Diagnostic Trace Service. (2) Select the Enable trace check box, and enable tracing for the MDM groups. (3) Restart the server and run the application again. (4) Check the errorTrace.log file to determine which problems were encountered.

HWN021516E The LSS specified *LSS name* on subsystem *Name of subsystem* is already at the maximum volume number (255). Volume creation can not be done on this LSS, please select a different one.

Explanation

Only 255 volumes can be created on a particular LSS. IBM Spectrum Control believes that this LSS is full because the highest numbered volume on the LSS is 255.

Action

If volume number 255 has been deleted, try rerunning a probe of the subsystem. Otherwise, either choose a different LSS or else delete volume #255 (and any others as necessary depending on the number of new volumes needed)

HWN021517E The connection to SMI-S provider for storage system *VPD of the storage system* could not be made.

Explanation

Creation of the CIM client failed.

Action

Ensure that the SMI-S provider is running and that the correct user ID and password have been set for this SMI-S provider in the IBM Spectrum Control UI.

HWN021520E The attribute *Name of the attribute* was not found.

Explanation

A method was called with wrong attributes.

Action

Try the same operation again. If the problem persists, enable high level tracing as explained in the Installation and Configuration Guide and contact IBM service.

Related reference

- [Getting support](#)

HWN021522E Host port *The WWPN of the host port* not assigned to Volume *The PK of the volume*

Explanation

The host port is not assigned to the volume.

Action

Specify a port that is assigned to the volume.

HWN021524E Indexed Properties *Names* don't match

Explanation

Indexed input parameters are not of the same size

Action

Make sure that all the indexed input parameters are of the same size.

HWN021529E An SMI-S provider has reported unexpected values: *IP and port of SMI-S provider.*

Explanation

The values returned by the SMI-S provider might indicate an SMI-S provider malfunction.

Action

If not already enabled, enable tracing for the Device server. To enable tracing for the Device server, go to the IBM Spectrum Control Information Center and search for "Configuring tracing". Restart the Device server, run the application again, and check the Device server log files for more information about this error.

HWN021530E The Volume - Port mapping can not be created. There are existing mappings that prevent this combination. *VolumeCOP: The ID of the volume , Port: The WWPN of the port that should be mapped to the volume*

Explanation

For FASTT it is not possible to merge existing mappings. That means, if you have volume A mapped to port 1, and volume B mapped to port 2, you will neither be able to map port 1 to volume B, nor port 2 to volume A.

Action

Build up mappings beginning with one volume-port mapping, and then add further volumes and ports. For Example, first map volume A and port 1, then volume A and port 2, finally volume B and port1 or port 2. The volume will be mapped to both ports. See also documentation for details.

HWN021531E SMI-S provider *The IP and port of the SMI-S provider* can not reach storage system *The VPD of the storage system*

Explanation

The SMI-S provider can not reach the subsystem that it manages.

Action

Check logs on SMI-S provider side to determine why it can't reach the device.

HWN021535E There is not enough space left in the storage pool *The primary key of the Poolon storage system The VPD of the storage system* to create a volume of *The requested volume size bytes.*

Explanation

The size of the volume to be created is too large for the selected storage pool.

Action

Choose or create a storage pool with enough space or create a smaller volume.

HWN021536E The CIM method *The CIM method that is not supported.* is not supported on the storage system *The VPD of the storage system*

Explanation

The storage system does not support the specified function.

Action

No action is required.

HWN021537E Could not create connection to SMI-S provider *The IP and port for the SMI-S provider.* Reason: *The exception returned by the SMI-S provider.*

Explanation

No connection to the SMI-S provider could be established.

Action

Check the given reason.

HWN021538E The username *The username that was used to connect to the SMI-S provider.* or password is wrong on SMI-S provider *The IP and port for the SMI-S provider.*

Explanation

The username and/or password specified for the SMI-S provider are not valid.

Action

Set correct username/password combination for the SMI-S provider.

HWN021539E The SVC with IP *The IP of the SVC.* which is managed by SMI-S provider *The IP and port for the SMI-S provider.* can not be discovered. The status is *The status of the SVC.* .

Explanation

The SMI-S provider can not reach the SVC. The SVC is not added to the IBM Spectrum Control repository.

Action

There is a SVC which is disconnected and not reachable from the SMI-S provider. Please check the status of this SVC and run discovery again.

HWN021540E The invocation of CIM method *Name of method* failed on SMI-S provider *Name of SMI-S provider* . The return code is *Return code of method*. Details provided by the SMI-S provider : *Description of Returncode*

Explanation

The extrinsic CIM method that was invoked on the given SMI-S provider failed.

Action

Ensure that the correct input parameters for the CIM method have been used.

HWN021600W *Operation Name of the operation. partially processed.*

Explanation

The request has been partially executed. Some operations as part of this request have failed. Please check the detailed error messages.

Action

No action is required.

HWN021601E *The operation(s) Operation_names failed.*

Explanation

Multiple operations failed. None of the requested operations was carried out.

Action

Check the error messages in the system log files for the jobs that failed.

HWN021602E *It is necessary to specify target ports for storage device VPD of the storage subsystem*

Explanation

This storage subsystem requires that target FCPorts are specified during the assignment operation.

Action

Re-run the operation and specify the target ports

HWN021603W *More storage volumes and ports than specified will loose access*

Explanation

The unassignment operation was executed successfully, but more storage volumes and ports have been unassigned than specified because the 'force' option was set to true for this command. All unassigned WWPNS and storage volumes are returned in the return object.

Action

Check that the additionally unassigned volumes/ports do not cause any problems.

HWN021604E *WWPNs and storage volumes to be unassigned not completely specified. Assigned WWPNS: All WWPNS that are assigned to the volumes in the host port collection , missing WWPNS: The WWPNS that are assigned but were not specified in the input parameter in the method unassign . Storage volumes to be unassigned not completely specified. Assigned storage volumes: Lists all storage volumes that are really assigned to the WWPNS. }, missing storage volumes: The storage volumes that are really*

assigned but were not specified in the input parameter in the method unassign

Explanation

Not all ports that belong to a host port collection were specified. In order to unassign a host port collection, all ports have to be specified. Not all storage volumes that are assigned to the WWPNS were specified in the input. In order to unassign the WWPNS, all storage volumes have to be specified.

Action

Specify all ports belonging to the host port collection. Specify all storage volumes that are assigned to the WWPNS. You can alternatively specify the force flag to automatically unassign all WWPNS and volumes that need to be included.

HWN021605I More storage volumes and ports than specified will gain access.

Explanation

The assignment operation was executed successfully, but more storage volumes and ports have been assigned than specified because the 'force' option was set to true for this command. All assigned WWPNS and storage volumes are returned in the return object.

Action

Check that the additionally assigned volumes/ports do not cause any problems.

HWN021606E WWPNS and storage volumes to be assigned not completely specified. Missing WWPNS: *The WWPNS that need to be assigned but were not specified in the input parameter.* . Storage volumes to be assigned not completely specified. Missing storage volumes: *The storage volumes that need to be assigned but were not specified in the input parameter.*

Explanation

Not all ports that belong to a host port collection were specified. In order to assign a host port collection, all ports have to be specified. Not all storage volumes that need to be assigned due to existing assignments to the WWPNS were specified in the input. In order to assign the WWPNS, all storage volumes have to be specified.

Action

Specify all ports belonging to the host port collection. Specify all storage volumes that need to be assigned to the WWPNS. You can alternatively specify the force flag to automatically assign all WWPNS and volumes that need to be included.

HWN021607E The client type *the client type* with description *the client description* is not supported on SMI-S provider *the SMI-S provider IP and port* for storage subsystem *the subsystem ID* of volumes *the volumeIDs of the subsystem* which were passed in

Explanation

The SMI-S provider for the storage device does not support the chosen client type

Action

Specify the client type that is supported by the SMI-S provider. Because the supported client types also depend on the SMI-S provider version, you might need to change or upgrade the SMI-S provider.

HWN021608E The target port *the target port ID* does not belong to storage subsystem *the subsystem ID* of volumes *the volumeIDs* of the *subsystem which were passed in*

Explanation

The target ports specified do not belong to the subsystem of the input volumes

Action

Specify target ports of the same subsystem as the volumes or specify no target ports

HWN021609E There is not enough space left in the storage pool *The primary key of the Pool* on storage system *The VPD of the storage system* to create *The number of volumes to create* volumes of *The total size needed bytes total*.

Explanation

The size of the volumes to be created is too large for the selected storage pool.

Action

Choose or create a storage pool with enough space or create smaller volumes.

HWN021610E The specified size *The size of the volume to create* is not supported on pool *The storage pool ID* Size has to be dividable by *Divisor* returned by *getSupportedSizeRange* and in between *Minimum* returned by *getSupportedSizeRange* and *Maximum* returned by *getSupportedSizeRange*

Explanation

This pool supports sizes within the given range and divisible by the given divisor only.

Action

Use a size which is divisible by the given divisor and within the provided range.

HWN021611E Volume *The volume ID* has mappings, it can not be deleted.

Explanation

A volume can not be deleted as long as it has mappings

Action

Delete all mappings of the volume.

HWN021612E The mapping between volume *The volume ID* and port *The initiator port wwpn* exists already

Explanation

A mapping that exists already can not be created again

Action

None

HWN021613E *The WWPN The WWPN not found can not be found on subsystem The subsystem*

Explanation

The subsystem does not know the WWPN given.

Action

The information in the database repository might be out of sync with the monitored devices in your environment. Rerun data collection jobs against those monitored devices to refresh the information stored in the database repository.

HWN021614E *The WWPNs The WWPNs without mappings have no mappings on storage system The storage system*

Explanation

The WWPNs do not have any mappings on this storage system.

Action

The information in the database repository might be out of sync with the monitored devices in your environment. Rerun data collection jobs against those monitored devices to refresh the information stored in the database repository.

HWN021615E *WWPNs WWPNs that can not share mappings can not share mappings on storage system Storage system}. There are existing mappings that prevent this.*

Explanation

For some devices (e.g. FASTT, HDS) it is not possible to merge existing mappings. That means, if you have volume A mapped to port 1, and volume B mapped to port 2, you will not be able to create any mapping that has both port1 and port2 included.

Action

Create all mappings at once; specify all ports and volumes to be mapped together. OR: Build up mappings beginning with one volume-port mapping, and then add further volumes and ports. For Example, first map volume A and port 1, then volume A and port 2, finally volume B and port1 or port 2. The volume will be mapped to both ports. See also documentation for details.

HWN021616E *Volumes VolumeIDs can not share mappings on storage system Storage system }. There are existing mappings that prevent this.*

Explanation

For some devices (e.g. FASTT, HDS) it is not possible to merge existing mappings. That means, if you have volume A mapped to port 1, and volume B mapped to port 2, you will not be able to create any mapping that has both volume A and volume B included.

Action

Create all mappings at once; specify all ports and volumes to be mapped together. OR: Build up mappings beginning with one volume-port mapping, and then add further volumes and ports. For Example, first map volume A and port 1, then volume A and port 2, finally volume B and port1 or port 2. The volume will be mapped to both ports. See also documentation for details.

HWN021617E The stored data for storage system *The storage system* is not in sync with the environment. Rerun data collection.

Explanation

The data stored for the storage system does not seem to be in sync with the environment

Action

Rerun data collection.

HWN021618E Modifying target ports is not supported by subsystem *the subsystem* .

Explanation

The subsystem does not support modifying target ports of existing mappings.

Action

Remove the mapping and re-create with the new set of target ports.

HWN021619E Modifying the target ports for mapping of initiator port *initiator port WWPN* and volume *volume name* will also modify the target ports of the following mappings: *port - volume list*

Explanation

The initiator port has mappings to more volumes than were specified. The mappings to all volumes will be modified.

Action

Specify all volumes that are impacted.

HWN021620I Modifying the target ports for mapping of initiator port *initiator port WWPN* and volume *volume name* will modify the target ports of more mappings than specified.

Explanation

The initiator port has mappings to more volumes than were specified. The mappings to all volumes will be modified.

Action

Check the job log to see the additionally modified mappings.

HWN021621E It is not supported to modify the target ports of existing mappings and create new mappings in one step. Modify the existing mappings first and then create the new mappings. Existing mappings: *port - volume list*

Explanation

It is not supported to modify the target ports of existing mappings and create new mappings in one step.

Action

Modify the existing mappings first and then create the new mappings.

HWN021622I Started modification of the assignment of volume *VolumeID* on subsystem *Subsystem* to initiator port *WWPN* . Target ports to add: *target ports to add* Target ports to remove: *target ports to remove*

Explanation

The task was started. The log will inform about the further process.

Action

No action required.

HWN021623I Finished modification of the assignment of volume *VolumeID* on subsystem *Subsystem* to initiator port *WWPN* . Target ports to add: *target ports added* Target ports to remove: *target ports removed*

Explanation

The task succeeded.

Action

No action required.

HWN021624E The modification of the assignment of volume *VolumeID* on subsystem *Subsystem* to initiator port *WWPN* failed. Target ports to add: *target ports to add* Target ports to remove: *target ports to remove*

Explanation

The task failed.

Action

Check the log for failure reason.

HWN021650E A timeout occurred while connecting to SMI-S provider *SMI-S provider IP and port*.

Explanation

There was a timeout on an attempted connection to an SMI-S provider. The service is unavailable or not responding, or the network may be congested.

Action

Check that:

1. The SMI-S provider is running.
2. The SMI-S provider is able to answer IBM Spectrum Control requests without delay.
3. There is a good network connection between IBM Spectrum Control and the SMI-S provider.

If the SMI-S provider is running correctly and the network connection is fine, increase the timeout and try again.

HWN021651E Job on SMI-S provider *SMI-S provider IP and Port in format IP:Port* failed. Job Status: *Job status* . Error code is *Error code* , error description: *Error description* . Check IBM Spectrum Control and SMI-S provider logs.

Explanation

An asynchronous job on the SMI-S provider failed.

Action

Check IBM Spectrum Control and SMI-S provider logs.

HWN021652E The process has timed out. Check the IBM Spectrum Control log files for more information.

Explanation

The process has taken longer than the specified timeout period. The process might be continuing even though the timeout has occurred.

Action

Check the log file at TPC_installation_directory\device\log\dmSvcTrace.log. If there are no additional error messages, then the process completed successfully in a longer amount of time than expected.

HWN021653E The attribute *Name of the attribute* was not found.

Explanation

A IBM Spectrum Control execution error occurred. Check IBM Spectrum Control logs.

Action

Check IBM Spectrum Control logs.

HWN021654E Pool ID *was not found*.

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Rerun the discovery process, and then run a data collection task.

HWN021655E Volume ID *The ID of the volume* was not found.

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Rerun the discovery process, and then run a data collection task.

HWN021656E Port ID *The ID of the port was not found.*

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Rerun the discovery process, and then run a data collection task.

HWN021657E Subsystem ID *The ID of the subsystem was not found.*

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Rerun the discovery process, and then run a data collection task.

HWN021658E Managed Disk ID *The ID of the MDisk was not found.*

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Rerun the discovery process, and then run a data collection task.

HWN021659E SMI-S provider *The ID of the SMI-S provider was not found*

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Rerun the discovery process, and then run a data collection task.

HWN021660E IO Group *The SVC IO Group was not found.*

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Rerun the discovery process, and then run a data collection task.

HWN021661E Extent *The storage extent external key was not found.*

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Rerun the discovery process, and then run a data collection task.

HWN021662E Physical volume *The physical volume external key was not found.*

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Rerun the discovery process, and then run a data collection task.

HWN021670E The client type *the client type* with description *the client description* is not unique on SMI-S provider *the SMI-S provider IP and port* } for storage subsystem *the subsystem ID* of volumes *the volumeIDs of the subsystem which were passed in*

Explanation

This SMI-S provider for the storage device does not have several entries for the chosen client type.

Action

Specify client type and description that are unique for the SMI-S provider. The supported client types also depend on the SMI-S provider version. Maybe you will need to change or upgrade the SMI-S provider.

HWN021671I The storage system *The storage system* was deleted from the database

Explanation

This storage system was deleted from the database.

Action

No action is required.

HWN021672E The *storage system name* storage system was not removed because other monitoring actions are running on the device.

Explanation

The storage system is in use by other monitoring actions and cannot be removed from the database repository at this time. For example, a probe schedule is collecting data about the storage system.

Action

Wait for the probe schedule or other action to complete and try to remove the storage system again.

HWN021673E The probe job on SMI-S provider *SMI-S provider IP and Port in format IP:Port* did not complete within the time limit of *Microseconds* microseconds. The job is *Percent complete* percent complete. Check the SMI-S provider log for job status. Job

information: *JobCOP* . Run the probe job again after the current job has completed.

Explanation

The probe job was started on the SMI-S provider, but did not complete in the specified time limit. The time limit can be set by using the `setdscfg` CLI command to modify the `CIMJobContext.JobRetrievalRetry` and `CIMJobContext.JobRetrievalSleep` parameters in the Db2 table.

- `CIMJobContext.JobRetrievalRetry` defines the maximum number of retries to check a job for completion.
- `CIMJobContext.JobRetrievalSleep` defines the wait time in microseconds between two retries.

Action

Check the SMI-S provider logs for status of the probe job. Use the `setdscfg` CLI command to change the CIM job retrieval parameters if necessary. Run the probe again to get the most recent data from the subsystem.

For more information about CLI commands, see the Command-line interface section of the Knowledge Center.

HWN021674E Job on SMI-S provider *SMI-S provider IP and Port* in format *IP:Port* returned unexpected results. Job information: *JobCOP* Job status: *JobState* , status description: *JobStatus* Check SMI-S provider log. Redo probe if the job completed.

Explanation

The job was started on the SMI-S provider, but returned an unexpected status.

Action

Check the SMI-S provider logs for job information. If the job completed, probe the subsystem again to get the up to date data.

HWN021675I Started creation of volume with size *Size* in pool *Pool* on subsystem *Subsystem*

Explanation

The task was started. The log will inform about the further process.

Action

No action required.

HWN021676I Volume creation completed successfully. New volume *VolumeID* created with size *Size* in pool *Pool* on subsystem *Subsystem* .

Explanation

The task succeeded.

Action

No action required.

HWN021677E Volume creation failed. The volume of size *Size* in pool *Pool* on subsystem *Subsystem* could not be created.

Explanation

The task failed.

Action

Check the log for failure reason.

HWN021678I Started assignment of volume *VolumeID* on subsystem *Subsystem* to initiator port *WWPN* .

Explanation

The task was started. The log will inform about the further process.

Action

No action required.

HWN021679I Finished assignment of volume *VolumeID* on subsystem *Subsystem* to initiator port *WWPN* .

Explanation

The task succeeded.

Action

No action required.

HWN021680E The assignment of volume *VolumeID* on subsystem *Subsystem* to initiator port *WWPN* failed.

Explanation

The task failed.

Action

Check the log for failure reason.

HWN021681I Started unassignment of volume *VolumeID* on subsystem *Subsystem* to initiator port *WWPN* .

Explanation

The task was started. The log will inform about the further process.

Action

No action required.

HWN021682I Finished unassignment of volume *VolumeID* on subsystem *Subsystem* to initiator port *WWPN* .

Explanation

The task succeeded.

Action

No action required.

HWN021683E The unassignment of volume *VolumeID* on subsystem *Subsystem* to initiator port *WWPN* failed.

Explanation

The task failed.

Action

Check the log for failure reason.

HWN021684I Started deletion of volume *VolumeID* on subsystem *Subsystem* .

Explanation

The task was started. The log will inform about the further process.

Action

No action required.

HWN021685I Volume deletion completed successfully. Volume *VolumeID* on subsystem *Subsystem* was deleted.

Explanation

The task succeeded.

Action

No action required.

HWN021686E Volume deletion failed. Volume *VolumeID* on subsystem *Subsystem* could not be deleted.

Explanation

The task failed.

Action

Check the log for failure reason.

HWN021687I Started modification of Pool *Pool display name* on subsystem *Subsystem display name* .

Explanation

The task was started. The log will inform about the further process.

Action

No action required.

HWN021688I Pool modification completed successfully. Pool *Pool display name* on subsystem *Subsystem display name* was modified.

Explanation

The task succeeded.

Action

No action required.

HWN021689E Pool modification failed. Pool *Pool display name* on subsystem *Subsystem display name* could not be modified.

Explanation

The task failed.

Action

Check the log for failure reason.

HWN021690I Started creation of *number volumes* volumes with size *Size* in pool *Pool* on subsystem *Subsystem*

Explanation

The task was started. The log will inform about the further process.

Action

No action required.

HWN021691I Created *number volumes* out of *total number volumes* volumes with size *Size* in pool *Pool* on subsystem *Subsystem*

Explanation

The task succeeded.

Action

No action required.

HWN021692E Volume creation failed. Created *number volumes* out of *total number volumes* volumes with size *Size* in pool *Pool* on subsystem *Subsystem*

Explanation

The task failed.

Action

Check the log for failure reason.

HWN021693W Warning: The task succeeded, but the database update failed. Run probe to update the database.

Explanation

The task succeeded, but the database update failed.

Action

Run probe for the subsystem to update the database.

HWN021700I Enumerating CIM Associator *The CIM association name which is being enumerated. for The name of the DB table which will be populated as result of this query.*

Explanation

The discovery or probe is currently enumerating a CIM associator. Inventory collection enumerates CIM classes in order to collect data for a particular IBM Spectrum Control entity, such as storage subsystem or storage volume.

Action

No action is required.

HWN021701I Enumerating CIM Class *The CIM class name which is being enumerated. for The name of the DB table which will be populated as result of this query.*

Explanation

The discovery or probe is currently enumerating a CIM class. Inventory collection enumerates CIM classes in order to collect data for a particular IBM Spectrum Control entity, such as storage subsystem or storage volume.

Action

No action is required.

HWN021702I Querying SMI-S provider

Explanation

An SMI-S provider query is in process

Action

No action is required.

HWN021703I Task starting on SMI-S provider *Identifier of the SMI-S provider..*

Explanation

The task is starting on the specified SMI-S provider.

Action

No action is required.

HWN021708I Initializing Collection for storage system *storage system identification*.

Explanation

Probe is being initialized.

Action

No Action is required

HWN021709I Collection for storage system *storage system identification* completed.

Explanation

Probe is completed.

Action

No action is required.

HWN021710I Discovering devices for SAN Volume Controller *The VPD of the SAN Volume Controller*.

Explanation

Prior to the discovery or probe, a fiber channel discovery on this SAN Volume Controller is issued.

Action

No action is required.

HWN021711I Discovery devices for SAN Volume Controller *The VPD of the SAN Volume Controller*. failed with error message *The exception which has occurred*.

Explanation

Fiber channel discovery issued prior to discovery or probe on the SAN Volume Controller was failing.

Action

Collection will continue without discovery.

HWN021712I Collecting Nodes for storage system *storage system identification*.

Explanation

The probe is currently traversing and storing CIM information that are related to the Nodes of the storage system.

Action

No action is required.

HWN021713I Collecting fibre channel ports for storage system *storage system identification.*

Explanation

The probe is currently traversing and storing CIM information that are related to the fibre channel ports of the storage system.

Action

No action is required.

HWN021714I Collecting volumes for storage system *storage system identification.*

Explanation

The probe is currently traversing and storing CIM information that are related to the volumes of the storage system.

Action

No action is required.

HWN021715I Traversing host to volume assignments for storage system *storage system identification.*

Explanation

The probe is currently traversing and storing CIM information that are related to host to volume assignment.

Action

No action is required.

HWN021716I Collecting pools and volumes for storage system *storage system identification.*

Explanation

The probe is currently traversing and storing CIM information that are related to the pools and volumes of the storage system.

Action

No action is required.

HWN021717I Collecting volume settings for storage system *storage system identification.*

Explanation

The probe is currently traversing and storing CIM information that are related to the volume settings of the storage system.

Action

No action is required.

HWN021718I Collecting client setting data for storage system *storage system identification.*

Explanation

The probe is currently traversing and storing CIM information that are related to the client setting data of the storage system.

Action

No action is required.

HWN021719I Perform collection post process tasks for storage system *storage system identification.*

Explanation

CIM based collection is completed and post collection tasks for the probe are performed.

Action

No action is required.

HWN021720I Flash enclosure is missing drive *flash_drive_identifier.*

Explanation

The drive is missing from its enclosure. Typically, all drives of a flash enclosure are installed.

Action

No action is required. Consider replacing the missing drive. However, the missing drive does not adversely affect the operation of the storage system or the ability to monitor the storage system.

HWN021724W SMI-S provider *SMI-S provider identifier* manages device(s) of type *device_type* which is supported through the native device interface or SNMP only.

Explanation

The specified device type is not supported through the SMI-S provider.

Action

Add the device again using a supported interface such as the CLI, the native API, or SNMP.

HWN021725I IBM Spectrum Control discovered/rediscovered a device with name *Identifier of the device.* on SMI-S provider *Identifier of the SMI-S provider..*

Explanation

The discovery found a device on an SMI-S provider it is examining.

Action

No action is required.

HWN021726I IBM Spectrum Control discovered/rediscovered no device on SMI-S provider *Identifier of the SMI-S provider..*

Explanation

The discovery found no device on an SMI-S provider.

Action

Check if this is correct that there is no device on the SMI-S provider.

HWN021727I IBM Spectrum Control discovery starting on SMI-S provider *Identifier of the SMI-S provider..*

Explanation

The discovery is starting on an SMI-S provider.

Action

No action is required.

HWN021728I IBM Spectrum Control discovery on SMI-S provider *Identifier of the SMI-S provider. is complete.*

Explanation

The discovery is finished on an SMI-S provider.

Action

No action is required.

HWN021729W IBM Spectrum Control discovery of Device type value is not supported.

Explanation

The device type specified is not supported.

Action

No action is required.

HWN021730W IBM Spectrum Control discovery of device value with code level value is not supported on SMI-S provider *Identifier of the SMI-S provider..*

Explanation

The code level specified is not supported.

Action

Upgrade the device to a version supported by the SMI-S provider.

HWN021731I Probing Volumes for Storage System: *value*.

Explanation

The probe is finding the volumes for this storage system.

Action

No action is required.

HWN021732I Number of Volumes Found Currently: *value*. Continuing to Probe Volumes.

Explanation

The Probe is finding the Volumes. This status update is to inform how many volumes have been processed at this point during the probe.

Action

No action is required.

HWN021733I *value* Volumes Found.

Explanation

This status message is to inform you of the total number of volumes found for this storage system or storage pool.

Action

No action is required.

HWN021734I Probing Disks for Storage System: *value*.

Explanation

The probe is finding the disks for this storage system.

Action

No action is required.

HWN021735I Number of Disks Found Currently: *value*. Continuing to Probe Disks.

Explanation

The probe is finding the disks. This status update is to inform how many disks have been processed at this point during the probe.

Action

No action is required.

HWN021736I *value* Disks Found.

Explanation

This status message to inform you of the total number of disks found for this storage system or storage pool.

Action

No action is required.

HWN021737I Probing Virtual Disks for Cluster: *value*

Explanation

The probe is finding the virtual disks for this SAN Volume Controller cluster.

Action

No action is required.

HWN021738I Number of Virtual Disks currently found: *value*. Continuing to probe Virtual Disks.

Explanation

The probe is finding the virtual disks. This status update is to inform of the number of virtual disks that have been processed at this point during the probe.

Action

No action is required.

HWN021739I *value* Virtual Disks found.

Explanation

This is the total number of Virtual disks found on the corresponding SAN Volume Controller cluster.

Action

No action is required.

HWN021740I Probing Views of Host Initiator access to Volumes.

Explanation

The probe is finding the Host Initiator access to Volumes.

Action

Check logs for SQLExceptions logged for ServiceUtils.getConnection().

HWN021741I *value* Views Found.

Explanation

This status message to inform you of the total number of Views for Host Initiator access to Volumes that are found for this storage system.

Action

Check logs for SQLExceptions logged for ServiceUtils.getConnection().

HWN021742E The SMI-S provider *SMI-S provider URL* is not managing storage subsystems.

Explanation

The SMI-S provider is managing switches and no storage subsystems.

Action

Specify the correct SMI-S provider which manages the storage subsystems.

HWN021743E The SMI-S provider *SMI-S provider URL* is not managing switches.

Explanation

The SMI-S provider is managing storage subsystems only and no switches.

Action

Specify the correct SMI-S provider which manages switches.

HWN021744E Cannot connect to a resource because of an SSL certificate error. Troubleshooting information: http://www.ibm.com/support/docview.wss?uid=swg21976237

Explanation

This communication problem might be caused by an error with the SSL certificate on the resource.

Action

To learn about how to troubleshoot the problem, go to <http://www.ibm.com/support/docview.wss?uid=swg21976237>.

Related reference

- [Resolving security certificate errors in IBM Spectrum Control V5.2.9 and later](#)

HWN021745I Cannot connect to a resource because of an SSL certificate error. Troubleshooting information: http://www.ibm.com/support/docview.wss?uid=swg21976237. An alternate resource will be used.

Explanation

This communication problem might be caused by an error with the SSL certificate on the resource.

An alternate resource has been identified and will be used instead for the fabric or switch.

Action

To learn about how to troubleshoot the problem, go to <http://www.ibm.com/support/docview.wss?uid=swg21976237>.

Related reference

- [Resolving security certificate errors in IBM Spectrum Control V5.2.9 and later](#)

HWN021746W SMI-S provider *Identifier of the SMI-S provider.* manages Cisco device types through SNMP only.

Explanation

The device type specified is not supported through the SMI-S provider.

Action

Use the add switch wizard to add and configure switches of the specified device type by using the SNMP interface.

HWN021747E Unable to add the specified switch by using SNMP. The switch is a Brocade switch and can be added only by using an SMI agent.

Explanation

Brocade switches must be added by using an SMI agent rather than by SNMP.

Action

Restart the Add Switch wizard and select Brocade as the vendor, and complete the information for the SMI agent.

HWN021800E Failed to get a database connection.

Explanation

Failed to get a database connection. Null was returned instead.

Action

Check logs for SQLExceptions logged for ServiceUtils.getConnection().

HWN021801E The server failed to get SMI-S provider entity from database.

Explanation

The server failed to get SMI-S provider entity from database. Either cursor on table T_RES_REGISTERED_CIMOM or contained object was null.

Action

No action is required.

HWN021802E Experienced SQL problems while working with database: *The SQL error.*

Explanation

Failed to work with database. Received SQL error instead.

Action

No action is required.

HWN021803W The server did not get userid and or password for SMI-S provider *The Service URL of the SMI-S provider* from database.

Explanation

The server found userid and or password being null in database.

Action

No action is required.

HWN021804E The server failed to access slp attributes for SMI-S provider *The Service URL of the SMI-S provider* from database.

Explanation

The server failed to access slp attributes for this SMI-S provider. Either cursor on table T_RES_SLP_ATTRIBUTES or contained object was null.

Action

No action is required.

HWN021805E CIMOMManager failed to get a database mapper of type *The type of the database mapper*.

Explanation

The CIMOMManager failed to get a database mapper. Unable to persist data to database.

Action

No action is required.

HWN021806E CIMOMManager failed to get a valid mapper result from *The type of the database mapper*.

Explanation

The CIMOMManager failed to get a valid mapper result from a database mapper. Unable to persist data to database.

Action

No action is required.

HWN021807E CIMOMManager failed to get a proxy for calling slp discovery.

Explanation

The CIMOMManager failed to get a proxy for calling slp discovery. Unable to discover SMI-S providers.

Action

Restart DiscoveryService if not running.

HWN021808E The device cannot be contacted through any of the following SMI-S providers *The comma separated list of IP and port for the SMI-S providers..* Possible causes are that the SMI-S providers are not accessible or the device is disconnected from the SMI-S providers.

Explanation

No connection to the SMI-S provider could be established.

Action

Check the given reason.

HWN021809E The host for SMI-S provider *The service URL of the SMI-S providers.* was not resolvable in DNS.

Explanation

The hostname of the SMI-S provider could not be translated to an IP address.

Action

Check the hostname.

HWN021810E The service URL for SMI-S provider *The service URL of the SMI-S providers.* is not valid.

Explanation

The hostname of the SMI-S provider could not be translated to an IP address.

Action

Check the URL fragments.

HWN021811I The operational status for device *The ID of the device.* on SMI-S provider *The service URL of the SMI-S provider.* has this value *The operational status vector.* .

Explanation

The operational status for this device was retrieved from the specific SMI-S provider. The value is a string representation of an vector of integers.

Action

No action is required.

HWN021812E The operational status for device *The ID of the device.* on SMI-S provider *The service URL of the SMI-S provider.* could not

be retrieved because SMI-S provider is in status *The SMI-S provider connection status.* .

Explanation

The operational status for this device was not retrieved from the specific SMI-S provider. The SMI-S provider is in a state which does not allow retrieval of operational status for the device.

Action

Action depends on returned SMI-S provider connection status.

HWN021813E Fabric ID *The ID of the fabric was not found.*

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Rerun the discovery process, and then run a data collection task.

HWN021814E The device *device id* cannot be contacted through the SMI-S provider *SMI-S provider service URL*.

Explanation

The SMI-S provider is in a state which does not allow retrieval of operational status for the device or the device is disconnected from the SMI-S provider.

Action

Check that the device can be reached from the machine that the SMI-S provider is running on.

HWN021899E Switch *The wwn of the switch. has no associated Fabric.*

Explanation

The SMI agent that has reported the switch is not reporting a fabric for the switch. The discovered switch is not persisted, due to this failure.

Action

Try restarting the switch SMI agent. If this does not fix the problem, contact customer support for your switch vendor.

HWN021901E The virtual disk size cannot exceed *maximum size* when creating space efficient virtual disks.

Explanation

When creating space efficient virtual disks, the maximum virtual disk size cannot be higher than the specified value.

Action

Provide a correct size for the virtual disk.

HWN021902E Invalid grain size. Valid values are *valid values*.

Explanation

The specified grain size parameter is invalid.

Action

Pass a valid parameter value.

HWN021903E Authentication to *ip or name of host* failed. Please specify correct authentication information.

Explanation

Cannot authenticate with given authentication information.

Action

Make sure username and password are correct.

HWN021904E Connection to *IP address or name of host* failed with following operating system exception: *exception text* . Please make sure IP address is correct and machine is up and running. If this is a SVC V4 machine, it could be that its RAS interface is not up. If this is a SVC V5, make sure the SMI-S provider is up and running.

Explanation

Connection to the host failed.

Action

Make sure IP address is valid and machine is up.

HWN021905E Connection to *IP address or name of host* failed with following operating system exception: *exception text* .

Explanation

Connection to the host failed.

Action

Contact IBM support.

HWN021906E Failed to get native API entity from database.

Explanation

The NAPIManager failed to get native API entity or related information from database.

Action

Configure Subsystem connection again, run Discovery and Probe again for the failing device. Ensure that everything completes successful.

HWN021907E The IP address *The service URL of the SMI-S providers.* was not resolvable in DNS.

Explanation

The hostname could not be translated to an IP address.

Action

Check the hostname.

HWN021908E Failed to get a proxy for calling NAPI discovery.

Explanation

The NAPIManager failed to get a proxy for calling NAPI discovery. Unable to discover subsystems.

Action

Restart DiscoveryService if not running.

HWN021909E There are no IO Groups available for Virtual Disk creation.

Explanation

A virtual disk creation using system chosen parameters was attempted but there are no valid IO Groups available to choose from.

Action

Probe the San Volume Controller again. Check San Volume Controller configuration and ensure there is at least one IO Group that has nodes. Should the problem persist, contact IBM support.

HWN021910E Managed Disk ID *The ID of the MDisk is not in unmanaged mode and cannot be added to the specified managed-disk group.*

Explanation

The Managed Disk specified in the command is not in unmanaged mode. To be a candidate for a managed-disk group, a managed disk cannot be part of another managed group. It also cannot be either offline or excluded.

Action

Choose a different Managed Disk ID to add to the specified managed-disk group.

HWN021911E Another probe of storage subsystem *The Name+Nameformat of the storage subsystem is already in progress.*

Explanation

Another probe for the same subsystem was already started and is in progress, so the new probe cannot be started.

Action

Probe this subsystem only after the previous probe for it is finished.

HWN021912E Other probes of storage subsystems *The list of Name+Nameformat of the storage subsystems* are already in progress.

Explanation

Other probes for the same subsystems were already started and are in progress, so the new probes cannot be started.

Action

Probe these subsystems only after the previous probes for them are finished.

HWN021913E IBM Spectrum Control Device Server could not write to directory *The directory*.

Explanation

The IBM Spectrum Control Device Server is not able to write to the directory. Possible reasons could be insufficient disk space, missing access privileges, etc.

Action

Check the free disk space and ensure that the access permissions for the directory are set correctly. If the problem still occurs, please contact IBM support.

HWN021914E SSH key file *The SSH key file name* is still in use, so it cannot be deleted.

Explanation

The SSH key is still used by IBM Spectrum Control to manage one or more SVC devices, so it cannot be deleted.

Action

No action is required.

HWN021915E IBM Spectrum Control Device Server could not delete the file *The file*.

Explanation

The IBM Spectrum Control Device Server is not able to delete the file. A possible reason could be missing or wrong access privileges.

Action

Ensure that the access permissions for the file are set correctly. If the problem still occurs, please contact IBM support.

HWN021916E The storage subsystem *subsystem ID* is not configured for file level management.

Explanation

File level management information is not available for the storage subsystem.

Action

Verify the storage subsystem is configured for file level management. If the storage subsystem is configured for file level management, run a storage subsystem probe. Then try the operation again.

HWN021917E An invalid parameter *Name of the parameter which was invalid* was specified. The corresponding file system mount point does not exist.

Explanation

An invalid parameter was specified. The corresponding file system mount point does not exist.

Action

Specify a valid file system mount point and try the operation again.

HWN021919E The cluster ID *The ID of the cluster.* was not found.

Explanation

The cluster ID specified in the command was not found.

Action

Some of the IBM Spectrum Control components might be out of sync. Run a discovery, then run a data collection task, such as a subsystem probe. Try the command again.

HWN021920E The export ID *The ID of the export.* was not found.

Explanation

The export ID specified in the command was not found.

Action

Some of the IBM Spectrum Control components might be out of sync. Run a discovery, then run a data collection task, such as a subsystem probe. Try the command again.

HWN021921E The specified activity or protocol could not be used to change the export *The ID of the export..*

Explanation

A valid activity or protocol is required to change an export.

Action

Specify a valid activity or protocol to change the export.

HWN021922E The file system ID *file_system_ID* was not found.

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Run the discovery process again, and then run a data collection task.

HWN021923E Invalid parameter *Name of the parameter which was invalid.* File system does not exist.

Explanation

The file system system does not exist.

Action

Specify a file system.

HWN021924E The parameter *Name of the parameter which was invalid* is not a valid parameter.

Explanation

The NFS position option is only used by NFS add or NFS change option.

Action

In order to use the NFS position option, you must also specify either the NFS add option or the NFS change option.

HWN021925E The fileset ID *fileset_ID* was not found.

Explanation

No DB row was found that corresponds to the key that was passed in as an input parameter.

Action

IBM Spectrum Control components may be out of sync. Run the discovery process again, and then run a data collection task.

HWN021926E The WAN-cache source ID *WAN_cache_source_id* was not found.

Explanation

The WAN-cache source ID specified in the command was not found.

Action

Some IBM Spectrum Control components may be out of sync. Run a discovery, and then run a data collection task, such as a subsystem probe. Try the command again.

HWN021927E The WAN-cache ID *WAN_cache_source_id* was not found.

Explanation

The WAN-cache ID specified in the command was not found.

Action

Some IBM Spectrum Control components may be out of sync. Run a discovery, and then run a data collection task, such as a subsystem probe. Try the command again.

HWN023000I The Optimization Execution task has started.

Explanation

This message is for informational purposes. No further investigation is required.

Action

No action is required.

HWN023001E The task to optimize the volumes was not completed successfully.

Explanation

The action was not completed because unexpected errors occurred.

Action

Open the logs for the task and complete the following actions:

- Examine the log messages that were generated before the errors occurred.
- Determine the current state of the volumes.

Depending on the error, it might be necessary to access the storage virtualizer. If this is the case, complete the following actions:

- Carry out further investigations to determine the cause of the errors.
- Resolve the issues that caused the errors.

When the issues are resolved, probe the storage virtualizer, and then rerun the task to optimize the volumes. If the issues cannot be resolved, contact IBM Software Support.

Related reference

-  [Getting support](#)

HWN023002I The Optimization Execution task has completed.

Explanation

This message is for informational purposes. No further investigation is required.

Action

No action is required.

HWN023003I The Optimization Execution task retrieved *number* recommendations

Explanation

The Optimization Execution request will process the number of recommendations retrieved.

Action

No action is required.

HWN023004I The Optimization Automation request persisted recommendations to be processed.

Explanation

The Optimization request persisted the recommendations to be processed.

Action

No action is required.

HWN023005I The Optimization Execution task updated the status of *number* recommendations.

Explanation

The Optimization Execution task has updated the status of the specified number of recommendations.

Action

No action is required.

HWN023006I The Optimization Automation request begins processing *number* recommendations.

Explanation

The Optimization request is processing the specified number of recommendations.

Action

No action is required.

HWN023007W The recommendation being processed contains a virtual disk that is no longer detected.

Explanation

The recommendation list contains a recommendation request where the virtual disk is no longer detected by IBM Spectrum Control. The recommendation will be ignored.

Action

If the analysis that generated the recommendation was performed sometime in the past, it is possible that an entity that was part of the recommendation is no longer detected by TPC. It would therefore be recommended to resubmit the analysis job to generate a new set of recommendations that would not contain entities that no longer exist.

HWN023008W The recommendation for virtual disk *vdisk name* contains a source storage pool that is no longer detected.

Explanation

The recommendation list contains a recommendation request where the source storage pool is no longer detected by IBM Spectrum Control. The recommendation will be ignored.

Action

If the analysis that generated the recommendation was performed sometime in the past, it is possible that an entity that was part of the recommendation is no longer detected by TPC. It would therefore be recommended to resubmit the analysis job to generate a new set of recommendations that would not contain entities that no longer exist.

HWN023009W The recommendation for virtual disk *vdisk name* contains a target storage pool that is no longer detected.

Explanation

The recommendation list contains a recommendation request where the target storage pool is no longer detected by IBM Spectrum Control. The recommendation will be ignored.

Action

If the analysis that generated the recommendation was performed sometime in the past, it is possible that an entity that was part of the recommendation is no longer detected by TPC. It would therefore be recommended to resubmit the analysis job to generate a new set of recommendations that would not contain entities that no longer exist.

HWN023010I Virtual disk *vdisk name* was successfully migrated from storage pool *source pool name* to storage pool *target pool name*.

Explanation

The Recommendation processed successfully migrated the virtual disk to the target pool.

Action

No action is required.

HWN023011W The recommendation for virtual disk *vdisk name* contains a virtual disk that does not exist in the source storage pool *source pool name* or the target storage pool *target pool name*.

Explanation

The recommendation list contains a recommendation request where the virtual disk does not exist in either the source or target storage pool. The recommendation will be ignored.

Action

If the analysis that generated the recommendation was performed sometime in the past, it is possible that an entity that was part of the recommendation is no longer in a relationship once detected by IBM Spectrum Control. It would therefore be recommended to resubmit the analysis job to generate a new set of recommendations that would not contain entities that no longer exist.

HWN023012W The recommendation for virtual disk *vdisk name* contains a non-mirrored virtual disk that is now a mirrored virtual disk.

Explanation

The recommendation list contains a recommendation request where a non-mirrored virtual disk is now a mirrored virtual disk. The recommendation will be ignored.

Action

If the analysis that generated the recommendation was performed sometime in the past, it is possible that an entity that was part of the recommendation is no longer in a relationship once detected by IBM Spectrum Control. It would therefore be recommended to resubmit the analysis job to generate a new set of recommendations that would not contain entities that no longer exist.

HWN023013W The recommendation for virtual disk *vdisk name* contains a mirrored virtual disk that is now a non-mirrored virtual disk.

Explanation

The recommendation list contains a recommendation request where a mirrored virtual disk is now a non-mirrored virtual disk. The recommendation will be ignored.

Action

If the analysis that generated the recommendation was performed sometime in the past, it is possible that an entity that was part of the recommendation is no longer in a relationship once detected by IBM Spectrum Control. It would therefore be recommended to resubmit the analysis job to generate a new set of recommendations that would not contain entities that no longer exist.

HWN023014I The recommendation for virtual disk *vdisk name* requires more space on target pool *target pool name* to be processed.

Explanation

The Recommendation processed requires space to be freed by dependent migrations before the recommendation can be processed. Once the dependent migrations have been completed, the recommendation will be retried with the expectation that the recommendation will succeed since the space required was freed.

Action

No action is required.

HWN023015I Virtual disk *vdisk name* will now be migrated from storage pool *source pool name* to storage pool *target pool name*.

Explanation

The Recommendation processed will now perform the virtual disk migration.

Action

No action is required.

HWN023016I Successfully added virtual disk copy to virtual disk *vdisk name*.

Explanation

The Recommendation processed successfully added a virtual disk copy to the specified virtual disk.

Action

No action is required.

HWN023017I Synchronization for virtual disk *vdisk name* has completed *synchronization percent%* and requires about *seconds to complete* seconds to complete.

Explanation

The synchronization for the virtual disk is not complete and will need to take the specified amount of time to finish. Keep in mind that the time to complete is only an approximation and competing activity within the virtualizer may either hasten or slow the synchronization. Therefore, the time to complete adjusts as the synchronization of the virtual disk progresses.

Action

No action is required.

HWN023018I Synchronization for virtual disk *vdisk name* has completed.

Explanation

The synchronization for the virtual disk has completed.

Action

No action is required.

HWN023019I Successfully removed a virtual disk copy from virtual disk *vdisk name*.

Explanation

The Recommendation processed successfully removed a virtual disk copy of the specified virtual disk.

Action

No action is required.

HWN023020I Successfully changed the synchronization rate of virtual disk *vdisk name* to *syncrate%*.

Explanation

The Recommendation processed successfully changed the synchronization rate to the specified amount.

Action

No action is required.

HWN023021I Successfully changed the primary copy of virtual disk *vdisk name*.

Explanation

The Recommendation processed successfully changed the primary copy of the virtual disk.

Action

No action is required.

HWN023022E There is no space available on target pool *target pool name* to migrate the virtual disk *vdisk name*.

Explanation

The Recommendation for the virtual disk cannot be migrated due to no available space in the target storage pool.

Action

Make more space available in the target storage pool or re-run the analysis so that a new recommendation with sufficient space can be accomplished.

HWN023023E Unable to submit request to add vdisk copy command for virtual disk *vdisk name* due to rc (*rc*).

Explanation

Could not submit the request to add a vdisk copy to the specified virtual disk due to an error with the request submission.

Action

Check subsequent messages for more details and additional log files that may surface communication errors under the device server log directory, otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN023024E Unable to complete request to add vdisk copy command for virtual disk *vdisk name* due to rc (*rc*).

Explanation

Could not complete the request to add a vdisk copy to the specified virtual disk due to an error with the request completion.

Action

Check subsequent messages for more details and additional log files that may surface communication errors under the device server log directory, otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN023025E Unable to submit request to get vdisk synchronization progress for virtual disk *vdisk name* due to rc (*rc*).

Explanation

Could not submit the request to get the vdisk synchronization progress of the specified virtual disk due to an error with the request submission.

Action

Check subsequent messages for more details and additional log files that may surface communication errors under the device server log directory, otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN023026E Unable to complete request to get vdisk synchronization progress for virtual disk *vdisk name* due to rc (*rc*).

Explanation

Could not complete the request to get the vdisk synchronization progress of the specified virtual disk due to an error with the request completion.

Action

Check subsequent messages for more details and additional log files that may surface communication errors under the device server log directory, otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN023027E Unable to submit request to remove vdisk copy command for virtual disk *vdisk name* due to rc (*rc*).

Explanation

Could not submit the request to remove a vdisk copy to the specified virtual disk due to an error with the request submission.

Action

Check subsequent messages for more details and additional log files that may surface communication errors under the device server log directory, otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN023028E Unable to complete request to remove vdisk copy command for virtual disk *vdisk name* due to rc (rc).

Explanation

Could not complete the request to remove a vdisk copy to the specified virtual disk due to an error with the request completion.

Action

Check subsequent messages for more details and additional log files that may surface communication errors under the device server log directory, otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN023029E Unable to submit request to change the synchronization rate for virtual disk *vdisk name* due to rc (rc).

Explanation

Could not submit the request to change the synchronization rate of the specified virtual disk due to an error with the request submission.

Action

Check subsequent messages for more details and additional log files that may surface communication errors under the device server log directory, otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN023030E Unable to complete request to change the synchronization rate for virtual disk *vdisk name* due to rc (rc).

Explanation

Could not complete the request to change the synchronization rate of the specified virtual disk due to an error with the request completion.

Action

Check subsequent messages for more details and additional log files that may surface communication errors under the device server log directory, otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN023031E Unable to submit request to change the primary copy for virtual disk *vdisk name* due to rc (rc).

Explanation

Could not submit the request to change the primary copy of the specified virtual disk due to an error with the request submission.

Action

Check subsequent messages for more details and additional log files that may surface communication errors under the device server log directory, otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN023032E Unable to complete request to change the primary copy for virtual disk *vdisk name* due to rc (*rc*).

Explanation

Could not complete the request to change the primary copy of the specified virtual disk due to an error with the request completion.

Action

Check subsequent messages for more details and additional log files that may surface communication errors under the device server log directory, otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN023033E The request failed. Message from failed request: *message*.

Explanation

The request failed due to an error processing the request.

Action

Contact IBM customer technical support if unable to determine the error from the request error message or the log files under the device server log directory.

Related reference

- [Getting support](#)

HWN023034E The Optimization Automation job completed with errors in the recommendations.

Explanation

The Optimization Automation job failed due to errors with the recommendations.

Action

Investigate the reason for the recommendation errors. This can be done by looking at the messages within the job log and taking the appropriate action where needed.

HWN023035W The Optimization Execution task completed with warnings.

Explanation

The Optimization Execution task encountered errors that were ignored during the processing of the recommendations.

Action

Investigate the reason for the recommendation errors. This can be done by looking at the messages within the job log and taking the appropriate action where needed.

HWN023036E The request failed because there were not enough extents in the storage pool.

Explanation

The request failed due to there not being enough space to satisfy the request.

Action

Make more space available in the target storage pool or re-run the analysis so that a new recommendation with sufficient space can be accomplished.

HWN023037E The request failed because the number of copies of this volume would exceed the limit.

Explanation

The request failed due to there being the expected limit of volume copies for the volume. It is possible that the volume has already been migrated to a new target pool.

Action

If the analysis that generated the recommendation was performed sometime in the past, it is possible that an entity that was part of the recommendation is no longer detected by TPC. It would therefore be recommended to resubmit the analysis job to generate a new set of recommendations that would not contain entities that no longer exist.

HWN023038E The request failed because the copy specified does not exist.

Explanation

The request failed due to the copy specified no longer existing. It is possible that the volume has already been migrated to a new target pool.

Action

If the analysis that generated the recommendation was performed sometime in the past, it is possible that an entity that was part of the recommendation is no longer detected by TPC. It would therefore be recommended to resubmit the analysis job to generate a new set of recommendations that would not contain entities that no longer exist.

HWN023039E The following exception occurred during a migration request: *exception*

Explanation

The migration request failed due to an exception The exception is caused by an internal error in the program.

Action

Contact IBM support with all the service log files so that the exception can be investigated further.

Related reference

- [Getting support](#)

HWN023040E The migration request for volume *vdisk name* is already being processed.

Explanation

The migration request failed for the specified volume because it is already being processed. This could be due to a concurrent job request to perform different migrations on the same volume. The first migration request for the volume will be preferred over any subsequent migration request for the same volume.

Action

Check the job logs for competing migration requests and if the wrong migration was performed, perform the analysis again and resubmit the optimization automation.

HWN023041W The request to migrate the mirrored volume *vdisk name* is suspended because the secondary volume is offline.

Explanation

To complete the migration request for the mirrored volume, the status of the secondary volume must change from offline to online. The status of the secondary volume is checked at regular intervals. If the status of the secondary volume changes to online, the migration request is resumed.

Action

If the migration request remains suspended, you can either cancel the execution of the recommendation, or open the Storage System GUI to bring the secondary volume back online.

HWN023042E The secondary copy needed for migration does not exist.

Explanation

The secondary copy is needed for volume migration since it is the means by which migration is achieved. Without the secondary copy, the migration cannot complete.

Action

Look at job logs to determine why the volume that was to be added was removed. If an action was taken that removed the new secondary copy before the copy was synchronized, then this error will occur. It is recommended to retry the recommendation and submit it separately from other recommendations.

HWN023043I The mirrored volume migration for volume *vdisk name* will be ignored.

Explanation

The mirrored volume migration for the specified volume will be ignored because the user specified that it should be ignored.

Action

No action is required.

HWN023044I The mirrored volume migration for volume *vdisk name* will result in the current secondary volume becoming the primary volume.

Explanation

The mirrored volume migration for the specified volume will result in the current secondary volume to become the primary volume and the new secondary volume will be migrated to the target pool. This is because the user specified this action.

Action

No action is required.

HWN023045I The mirrored volume migration for volume *vdisk name* will result in the primary volume being migrated to the target pool.

Explanation

The mirrored volume migration for the specified volume will result in the primary volume to be migrated to the target pool because the user specified this action.

Action

No action is required.

HWN023046I The Migration of the previously abandoned Optimization Automation job has started.

Explanation

If an optimization automation job was in progress prior to a server shutdown, the job will resume if any migrations are pending. This message is for informational purposes. No further investigation is required.

Action

No action is required.

HWN023047I The Migration of the previously abandoned Optimization Automation job has completed.

Explanation

This message is for informational purposes. No further investigation is required.

Action

No action is required.

HWN023048I The Optimization Automation cancellation job *jobname* has started.

Explanation

This message is for informational purposes. No further investigation is required.

Action

No action is required.

HWN023049E The Optimization Automation cancellation job completed with errors.

Explanation

The cancellation of the Optimization Automation job failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

-  [Getting support](#)

HWN023050I The Optimization Automation cancellation job *jobname* has completed.

Explanation

This message is for informational purposes. No further investigation is required.

Action

No action is required.

HWN023051I The Optimization Automation job *jobname* will be canceled.

Explanation

The request to cancel an optimization job will commence being canceled.

Action

The result of the cancellation can be located in the job log of the job that is being canceled.

HWN023052W The Optimization Automation job is not in progress.

Explanation

The request to cancel an optimization job could not be canceled because it is not being processed.

Action

Check to see if the optimization job that was to be canceled has completed.

HWN023053I The migration of volume *vdisk name* has been canceled.

Explanation

The migration of the volume was canceled by the user. No further investigation is required.

Action

No action is required.

HWN023054W The Optimization Automation job was canceled.

Explanation

The Optimization Automation job was canceled by the user.

Action

Check if the cancellation was valid, and if so, no action is necessary. If not, look at the cancel request to determine who initiated the cancellation..

HWN023055I The volume that was chosen for transformation, *vdisk name*, is a secondary volume in a mirrored volume relationship. The secondary volume will be migrated to the specified target pool or converted as specified.

Explanation

The secondary volume will be converted or moved as specified.

Action

No action is required.

HWN024000I An optimization analysis task was started.

Explanation

This message is for informational purposes only.

Action

None.

HWN024001I The analysis is completed.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024002W Unable to retrieve any policy for Tier value.

Explanation

No policy was found for this tier. The input storage pool might have moved to another tier.

Action

If there have been changes to the storage pool tier since the analysis job was created, create a new analysis job.

HWN024003I Analyzed *number_of_volumes* volumes on tier *tier_number* for storage virtualizer *subsystem_name*.

Explanation

This message is for informational purposes. No further investigation is required.

Action

No action is required.

HWN024006W No target pools in *subsystem value* were selected.

Explanation

Source pools in the specified subsystem were selected, but no target pools in that subsystem were selected. IBM Spectrum Control will only move volumes across pools in the same subsystem.

Action

Select target pools in the specified subsystem and retry.

HWN024011W Destination storage pool *value* in subsystem *value* was not considered. Reason: *value*.

Explanation

This message is for informational purposes.

Action

None.

HWN024012I It is recommended that *number_of_volumes* volumes on tier *source_tier_number* are moved to tier *target_tier_number*.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024015I The optimization analysis of the *value* subsystem was started.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024016W Volume *value* is already in the destination storage pool *value*. No recommendations will be generated for the volume.

Explanation

The volume is already in the storage pool that was selected as the destination. No recommendations will be generated for the volume.

Action

Select a destination storage pool that is not the same as the source storage pool.

HWN024018W No destination storage pools in Tier *value* have been specified for subsystem *value*.

Explanation

Optimization Analysis cannot balance storage pools when no destination pools have been specified.

Action

Specify destination storage pools for the balance operation.

HWN024019W The following pools on tier *tier_number* on the *storage_system* storage system cannot be balanced by redistributing or re-tiering volumes: *pool_names*.

Explanation

The pools were not balanced either because volumes could not be re-tiered to balance the pools, or an insufficient number of volumes were re-tiered to reduce the activity level of the pools below the specified threshold.

Action

Complete one or more of the following actions, and then rerun the task to balance the pools:

- Select a wider range of pools on the tiers that you want to balance.
- Select a wider range of target pools on tiers that are higher and that are lower than the tiers that you want to balance.
- Increase the value for the activity limit of the pools that you want to balance.

HWN024020I Started analysis to balance pools on tier *value*.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024021W The *pool_name* pool on tier *tier_number* on the *storage_system* storage system cannot be balanced by redistributing the volumes.

Explanation

The pool was not balanced because an insufficient number of volumes could be redistributed across the pools to reduce the activity level of the pool below the specified threshold.

Action

Complete one or more of the following actions, and then rerun the task to balance the pools:

- Select a wider range of pools on the tiers that you want to balance.
- Increase the value for the activity limit of the pools that you want to balance.

Alternatively, run the Analyze Tiering wizard to automatically balance the pools. If the pools cannot be balanced across the tier, the most active volumes are re-tiered to balance the pools.

HWN024027I Storage Pool *pool name* has insufficient available space for volume *volume name* in storage pool *pool name*.

Explanation

The Optimization Balance Analysis determined that the specified volume could not be moved to the specified storage pool due to insufficient space. This does not mean that the specified storage pool is a good candidate in other aspects, just that the storage pool had insufficient space.

Action

If the storage pool activity values can not be reduced sufficiently by the Balance Analysis, consider resolving this issue by either modifying the storage pool or including additional storage pools in the Balance Analysis.

HWN024030W One or more entities specified as input for the analysis could not be found or pools or volumes in some input entities could not be found.

Explanation

The Optimization Analysis task was not able to find all of the input entities specified when the task was created or was not able to find pools or volumes in such entities. The entities may have been deleted or removed.

Action

If all input entities should still exist, check to see if IBM Spectrum Control reports any errors for these entities. If this situation is expected due to intentional environment changes, then no action is necessary.

HWN024031W One or more entities specified as candidate destinations for the analysis could not be found.

Explanation

The Optimization Analysis task was not able to find all of the candidate destination entities specified when the task was created. The entities may have been deleted or removed.

Action

If all candidate destination entities should still exist, check to see if IBM Spectrum Control reports any errors for these entities. If this situation is expected due to intentional environment changes, then no action is necessary.

HWN024032W For one or more mirrored volumes, both the primary and the secondary volume copies were chosen for transformation. You cannot transform both volume copies in the same transform task. Only the primary volume copies are included for transformation. You can transform the secondary volume copies in a separate transformation.

Explanation

If you are transforming both the primary and secondary copies of a mirrored volume, you must include the volume copies in separate transformations.

Action

On the Volumes page, choose the secondary volume copies that were excluded from this transformation and run the transform storage analysis for those volumes.

HWN024033W The volume *volume name* cannot be analyzed because it is not in a capacity pool.

Explanation

When you restrict optimization to target pools in the same capacity pool, the volumes that you select for analysis must belong to a capacity pool.

Action

Add the volume to a capacity pool and run the Analyze Tiering wizard again. Alternatively, choose another option for selecting target pools such as manual selection.

HWN024034W The pool *pool name* cannot be analyzed because the pool is not in a capacity pool.

Explanation

When you restrict optimization to target pools in the same capacity pool, the pools that you select for analysis must belong to a capacity pool.

Action

Add the pool to a capacity pool and run the Analyze Tiering wizard again. Alternatively, choose another option for selecting target pools such as manual selection.

HWN024035W The storage virtualizer *system name* cannot be analyzed because the storage virtualizer is not in a capacity pool.

Explanation

When you restrict optimization to target pools in the same capacity pool, the storage virtualizer that you select for analysis must belong to a capacity pool.

Action

Add the storage virtualizer to a capacity pool and run the Analyze Tiering wizard again. Alternatively, choose another option for selecting target pools such as manual selection.

HWN024036W The operation to transform the volumes on the *subsystem name* storage virtualizer cannot be completed because the destination pools were not available.

Explanation

The destination pools that were selected to transform the volumes could not be detected.

Action

Run the Transform Storage wizard again.

HWN024037E An unexpected error occurred. The operation to transform the volumes on the *subsystem name* storage virtualizer cannot be completed because the destination pools were not identified.

Explanation

The pools that were selected to transform the volumes could not be identified.

Action

If the problem persists, use the Service tool to collect trace data and send it to IBM Software Support.

Related reference

- [Getting support](#)

HWN024043I The capacity pools of the source volumes were selected as the target pools.

Explanation

In the Analyze Tiering wizard, the option to restrict the selection of target pools to target pools in the same capacity pool was selected.

Action

No action is required.

HWN024046I The option that was selected to handle volumes with mirrored volumes is: After optimization, set the copy of the secondary volume in the destination pool as the primary volume. The original secondary volume remains the secondary volume.

Explanation

The option that was selected by the user to include or exclude volumes with mirrored volumes in the analysis.

Action

No action is required.

HWN024047I The number of days for collecting performance data to analyze the volumes is set to *performance_data_collection_period*.

Explanation

Only the performance data that was collected within the specified period is used to analyze the volumes.

Action

No action is required.

HWN024050I Automatic tiering was selected to tier the volumes.

Explanation

The default option, automatic tiering, was selected by the user in the Analyze Tiering wizard.

Action

No action is required.

HWN024051I The tiering analysis is starting.

Explanation

The source volumes that you selected are analyzed to determine whether they require re-tiering.

Action

No action is required.

HWN024052I Tier *tier#* has an I/O density threshold value of *value* per second per GiB.

Explanation

User-defined values for the I/O density thresholds were specified for the storage tiers.

Action

No action is required.

HWN024053I Tier *tier#*, has a file age threshold value of *value* percent of files last accessed within *time_unit*.

Explanation

User-defined threshold values were specified for file usage for the storage tier.

Action

No action is required.

HWN024054I The real capacity for the thin-provisioned volumes is set to *value unit*.

Explanation

The percentage or amount of the current capacity of the volume that is allocated when the volume is converted to a thin-provisioned volume was specified by the user in the Transform Storage wizard.

Action

No action is required.

HWN024055I The auto expand property of the thin-provisioned volumes is set to *yes/no*.

Explanation

Whether the auto expand property was enabled or disabled by the user in the Transform Storage wizard.

Action

No action is required.

HWN024056I The warning level for thin-provisioned volumes is set to *value %*.

Explanation

The warning level threshold that was specified by the user in the Transform Storage wizard for thin-provisioned volumes.

Action

No action is required.

HWN024057I The grain size that was specified for the thin-provisioned volumes is *grain_size* KiB.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024058I The real capacity for the compressed volumes is set to *value unit*.

Explanation

The percentage or amount of the current capacity of the volume that is allocated when the volume is converted to a compressed volume was specified by the user in the Transform Storage wizard.

Action

No action is required.

HWN024059I The auto expand property for the compressed volumes is set to *yes/no*.

Explanation

Whether the auto expand property was enabled or disabled by the user in the Transform Storage wizard.

Action

No action is required.

HWN024060I The warning level for the compressed volumes is set to *value*.

Explanation

The warning level threshold that was specified by the user in the Transform Storage wizard for compressed volumes.

Action

No action is required.

HWN024061I The option that was selected to handle volumes with mirrored volumes is: After optimization, set the secondary volume as the primary volume. The volume in the destination pool is the secondary volume.

Explanation

The option that was selected by the user to include or exclude volumes with mirrored volumes in the analysis.

Action

No action is required.

HWN024062I The option that was selected for mirrored volumes is: Do not optimize volumes with mirrored volumes.

Explanation

The option that was selected by the user to include or exclude volumes with mirrored volumes in the analysis.

Action

No action is required.

HWN024066I Tier *tier#* has an I/O rate threshold value of value I/O per second.

Explanation

User-defined values for the I/O rate thresholds were specified for the storage tiers. For each tier, the lower threshold value is reported, so the lowest tier is shown with a threshold of 0.0 I/O per second.

Action

No action is required.

HWN024067W Recommendations cannot be generated for *number_of_volumes* volumes because the volumes do not meet the tiering criteria for tier *current_tier_number* or for any lower tier.

Explanation

Based on the tiering criteria that was selected, one or more volumes should be moved to lower tiers. However, recommendations to move the volumes cannot be generated because the volumes do not match the tiering criteria that was selected for the lower tiers.

Action

No action is required.

HWN024068W Recommendations cannot be generated to move *number_of_volumes* volumes from *source_tier* to tier *target_tier_number* due to the pool activity limit value.

Explanation

One or more volumes could not be moved to the specified tier without causing one or more storage pools in the tier to be too active.

Action

No action is required.

HWN024069W Recommendations cannot be generated to move *number_of_volumes* volumes from tier *source_tier* to to tier *target_tier_number* because the destination storage pools do not have enough space.

Explanation

The volumes cannot be added to the destination storage pools that were selected because of insufficient available space.

Action

Reduce the allocated capacity in the destination pools or choose destination pools with sufficient available space when you run the Analyze Tiering wizard.

HWN024070I The analysis to optimize subsystem *storage_subsystem* was completed.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024071I The option that was selected was to restrict the placement of volumes in capacity pools to destination storage pools in the same capacity pool.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024072W No file age information for volume *volume name*.

Explanation

The file age threshold can not be applied to this volume as no file scan information exists.

Action

No action is required. To enable file age threshold analysis of this volume, run a file system scan for any filesystems contained on the volume and rerun the analysis.

HWN024073W Storage pool {0} in tier {1} needs at least one additional storage pool in the same tier for the Balance Analysis to run on this tier.

Explanation

Storage pools are balanced within their tier. An input storage pool can not be balanced if no other input storage pool is of the same tier.

Action

Ensure storage pool tier assignments are correct and that all necessary storage pools are selected as input for the Balance Analysis.

HWN024074W Storage pool {0} in tier {1} and capacity pool {2} needs at least one additional storage pool in the same tier and capacity pool for the Balance Analysis to run within this capacity pool and on this tier.

Explanation

Storage pools are balanced within their tier and capacity pool boundary. An input storage pool can not be balanced if no other input storage pool is of the same tier and capacity pool.

Action

Ensure storage pool tier and capacity pool assignments are correct and that all necessary storage pools are selected as input for the Balance Analysis.

HWN024075W *number_of_volumes* volumes from storage pool *pool* could not be moved to the destination storage pools because the destination storage pools do not have enough space.

Explanation

The volumes cannot be moved to the destination storage pools that were selected because of insufficient available space.

Action

Add more capacity to the destination storage pools or choose destination storage pools with sufficient available space. If there are other reasons why the volumes could not be moved to the destination storage pools, adding more capacity to the destination storage pools or choosing destination storage pools with sufficient available space may still not generate recommendations.

HWN024076W *number_of_volumes* volumes from storage pool *pool* could not be moved to the destination pools because the destination storage pools are not in the same capacity pool.

Explanation

The volumes cannot be moved to the destination storage pools that were selected because the destination storage pools are not in the same capacity pool as the source storage pool.

Action

Choose destination storage pools that are in the same capacity pool as the source storage pool. If there are other reasons why the volumes could not be moved to the destination storage pools, choosing destination storage pools in the same capacity pool as the source storage pool may still not generate recommendations.

HWN024077W *number_of_volumes* volumes from storage pool *pool* could not be moved to the destination storage pools because the destination storage pools would have exceeded the pool activity limit value.

Explanation

The volumes cannot be moved to the destination storage pools that were selected because moving the volumes would have caused the destination storage pools' activity to exceed the activity limit value.

Action

Choose destination storage pools with lower pool activity values. If there are other reasons why the volumes could not be moved to the destination storage pools, choosing destination storage pools with lower activity values may still not generate recommendations.

HWN024078W *number_of_volumes* volumes from storage pool *pool* could not be moved to the destination storage pools because the destination storage pools already have a volume copy.

Explanation

The volumes cannot be moved to the destination storage pools that were selected because a copy of the volume already exists on the destination storage pools.

Action

Choose destination storage pools with do not have volume copies. If there are other reasons why the volumes could not be moved to the destination storage pools, choosing destination storage pools without volume copies may still not generate recommendations.

HWN024079W Because of an internal error, the number of volumes in the *pool* storage pool that could not be moved to destination storage pools is *number_of_volumes*.

Explanation

The selected destination storage pools were not valid destinations for these volumes.

Action

Choose different destination storage pools.

HWN024080W Destination storage pool *pool* already contains a copy of storage volume *volume*.

Explanation

Moving the storage volume to the destination storage pool will result in both the storage volume and its copy being in the same storage pool.

Action

Choose a destination storage pool with no volume copies.

HWN024081W Because the destination storage pool does not have sufficient available space, the *volume* storage volume in the *source_pool* storage pool cannot be moved to the *destination_pool* destination storage pool.

Explanation

The destination storage pool that was selected does not have sufficient available space.

Action

Allocate more space to the destination storage pool or select a destination storage pool that has sufficient available space.

HWN024082W Because the destination storage pool contains a copy of the mirrored volume, the *volume* volume in the *source_pool* storage pool cannot be moved to the *destination_pool* destination storage pool.

Explanation

The primary volume and the secondary volume of a volume in a mirrored volume relationship cannot be placed in the same storage pool.

Action

Select a destination storage pool that does not contain a copy of the mirrored volume.

HWN024083W Because of an internal error, the *volume* storage volume in the *spool* storage pool could not be moved to the *destination_pool* destination storage pool.

Explanation

The storage pool is not a valid destination storage pool for the volume.

Action

Choose a different destination storage pool.

HWN024084W Because the destination storage pools contain one or more copies of the mirrored volumes, the number of volumes that could not be moved from tier *source_tier* to tier *target_tier* is *number_of_volumes*.

Explanation

The primary volume and the secondary volume of a volume in a mirrored volume relationship cannot be placed in the same storage pool.

Action

Ensure that one or more of the destination storage pools that are selected do not contain copies of the mirrored volumes.

HWN024085W The *pool_name* storage pool cannot be balanced because the tier level of the pool was reset to none.

Explanation

To balance storage pools, the tier level of each storage pool must be set.

Action

Set the tier level for the storage pool, and then rerun the task that was created to balance the pools.

HWN024086E Recommendations cannot be generated because the tier level of the *destination_pool_name* destination storage pool was reset to none.

Explanation

To analyze the tiering of the volumes, the tier level of both the source and the destination storage pools must be set.

Action

Set the tier level for the storage pool, and then rerun the task to analyze the tiering of the volumes.

HWN024087W Recommendations cannot be generated for one or more of the volumes because collocated volumes cannot be placed in the same destination storage pool.

Explanation

Based on the tiering criteria, the resources, and the options that were selected, all of the collocated volumes cannot be placed in the same destination pool.

Action

To determine whether further action can be taken, check the preceding messages in the log file. If, for example, one of the collocated volumes in the source pool was a mirrored volume, and the option to exclude mirrored volumes was selected, then recommendations to collocate the volumes are not generated. To resolve this particular issue, rerun the analysis and include mirrored volumes in the analysis.

HWN024088I The option to collocate volumes that are assigned to the same server or hypervisor was selected.

Explanation

If one or more of the volumes that are assigned to the same server or hypervisor and in the same source pool require optimization, then recommendations are only generated if all of the related volumes can be placed in the same destination pool.

Action

No action is required.

HWN024089I The option to collocate volumes that are assigned to the same server or hypervisor was not selected.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024090W Because the storage pools do not meet the service class requirements, the number of volumes that cannot be moved is *no_volumes*.

Explanation

Volumes that are assigned a service class must be placed in pools that meet the service class requirements of the volumes.

Action

Ensure that pools that meet the service class requirements of the volumes are selected when you run the Balance Pools analysis.

HWN024091W If the recommendation to move the *volume_name* volume to the *storage_pool_name* storage pool is implemented, the service class requirements of the *volume_name* volume cannot be met.

Explanation

Volumes that are assigned a service class should be placed in pools that meet the service class requirements of the volumes.

Action

Ensure that the destination pools that you select meet the service class requirements of the volumes in the source pools.

HWN024092W Recommendations cannot be generated to move *number_of_volumes* volumes from tier *source_tier* to tier *target_tier_number* because the destination storage pools do not meet the service class requirements of the volumes.

Explanation

The volumes cannot be added to the destination storage pools that were selected because the storage pools do not meet the service class requirements.

Action

Ensure that pools that meet the service class requirements of the volumes are selected.

HWN024093I The number of volumes on tier *tier_level* that were not analyzed because of the instruction to exclude mirrored volumes from the analysis is *number_of_volumes* volumes.

Explanation

The option to exclude mirrored volumes from the analysis was selected.

Action

No action is required.

HWN024094W Valid target pools were not selected for the *subsystem name* storage virtualizer.

Explanation

When volumes are optimized, they are moved to destination pools in the same storage virtualizer. The destination pool is chosen from the list of target pools that are selected when the analysis is run.

Action

Ensure that target pools for all of the resources that are being analyzed are selected.

HWN024095I The grain size for the thin-provisioned volumes was set to the default value of *grain_size* KiB.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024096W Volumes in the *pool_name* pool on tier *tier_level* cannot be moved to a higher tier to reduce the activity level of the pool to the user-defined level.

Explanation

No suitable target pools were found on higher tiers to relocate the most-active volumes.

Action

No action is required.

HWN024097W Volumes in the *pool_name* pool on tier *tier_level* cannot be moved to a lower tier to reduce the activity level of the pool to the user-defined level.

Explanation

No suitable target pools were found on lower tiers to relocate the most-active volumes.

Action

No action is required.

HWN024098W Cannot generate recommendations to tier volumes from the *storage_system_name* storage system because all of the source volumes are in the selected destination storage pools.

Explanation

To generate recommendations for tiering volumes, there must be at least one destination storage pool that does not include the volumes that are in the source storage pools.

Action

Select at least one destination storage pool that does not include the volumes that are in the source storage pools.

HWN024099I The number of volumes that were excluded from the analysis to plan the tiering of the *storage_system_name* storage system is *vols_count*. The volumes were excluded because performance data is not available for the volumes.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024100I The number of volumes that were excluded from the analysis to plan the tiering of the *storage_system_name* storage system is *vols_count*. The volumes were excluded from the analysis because the capacity of the volumes is zero.

Explanation

The volumes that were excluded from the analysis might be thin-provisioned volumes that are not being used and that have an initial capacity of zero.

Action

No action is required.

HWN024101I The number of volumes that were excluded from the analysis to plan the tiering of the *storage_system_name* storage system is *vols_count*. The volumes were excluded from the analysis because the volumes are not assigned to pools that are tiered or the thresholds were not defined for the tiers.

Explanation

To plan tiering, the volumes must belong to pools that are assigned a tier level with defined performance thresholds.

Action

Define the criteria for the tiers in your storage environment and ensure that the pools that you want to include in the analysis are assigned a tier level.

HWN024102W The recommendation to move the *storage_volume_name* volume from the *source_pool_name* storage pool to the *target_pool_name* storage pool was not generated because the status of the destination pool is offline or excluded.

Explanation

The recommendation was not generated and implemented because it would result in the volume becoming offline

Action

Select another destination storage pool or run a probe of the system that contains the destination storage pool. The status of the storage pool might have changed since the last probe.

HWN024103I Reclaiming volumes

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024104I Planning for tiering volumes

Explanation

This message is for informational purposes only.

Action

No action is required.

HWN024105W The recommendation to move the *storage_volume_name* volume from the *source_pool_name* storage pool to the *target_pool_name* storage pool will not be executed because the status of the destination pool is offline or excluded.

Explanation

The recommendation was not implemented because it would result in the volume becoming offline.

Action

No action is required.

HWN024106W The recommendation to move the *storage_volume_name* volume from the *source_pool_name* storage pool was not generated because the status of the volume is offline.

Explanation

An offline volume cannot be moved.

Action

Select another volume or run a probe of the system that contains the selected volume. The status of the volume might have changed since the last probe.

HWN024107W The recommendation to move the *storage_volume_name* volume from the *source_pool_name* storage pool to the *target_pool_name* storage pool will not be executed because the status of the volume is offline.

Explanation

An offline volume cannot be moved.

Action

No action is required.

HWN024108E The recommendations can't be shown because the analysis was not completed.

Explanation

The analysis to generate recommendations was not saved.

Action

Verify that the database is available and check the log files for the Data and the Device server. The default location of log files for the Data server is:

- Windows: TPC_installation_directory\data\log
- Linux and AIX: TPC_installation_directory/data/log

The default location of log files for the Device server is:

- Windows: TPC_installation_directory\device\log
- Linux and AIX: TPC_installation_directory/device/log

Related reference

-  [Getting support](#)

HWN024109W The data for the previous analysis of the *storage_subsystem* storage system was not deleted.

Explanation

Analysis data was not deleted.

Action

Verify that the database is available and check the log files for the Data and the Device server. The default location of log files for the Data server is:

- Windows: TPC_installation_directory\data\log
- Linux and AIX: TPC_installation_directory/data/log

The default location of log files for the Device server is:

- Windows: TPC_installation_directory\device\log
- Linux and AIX: TPC_installation_directory/device/log

HWN024110E Volumes reclamation analysis failed for *storage_subsystem* storage subsystem.

Explanation

Failed to work with database. Received SQL error instead.

Action

Ensure database server is up and running. Check device and data server logs to get error details.

HWN024111W Recommendations cannot be generated to move *number_of_volumes* volumes from tier *source_tier* to tier *target_tier_number* because there is no potential destination pool assigned to the recommended tier.

Explanation

There must be at least one potential destination pool assigned to the recommended tier in order to try a placement analysis.

Action

Select at least one pool assigned to this recommended tier, as potential destination, when you run the Analyze Tiering wizard.

HWN024112W Cannot generate recommendations to tier volumes from the *storage_system_name* storage system because the source storage pools and the selected destination storage pools are assigned to the same tier.

Explanation

To generate recommendations for tiering volumes, there must be at least one destination storage pool that is not assigned to the same tier as the source storage pools.

Action

Select at least one destination storage pool that does not belong to the same tier as the source storage pools.

HWN024200I The days of the week to include in the analysis: *days_of_week*.

Explanation

Only the performance data that was collected during the specified days of the week is used to analyze the volumes.

Action

No action is required.

HWN024201I The time window for the performance data to include in the analysis is set to *start time - end time*.

Explanation

Only the performance data that was collected during the specified time window is used to analyze the volumes.

Action

No action is required.

HWN024202I The time window for the performance data to include in the analysis is set to *start time - end time*. The end time occurs on the next day.

Explanation

Only the performance data that was collected during the specified time window is used to analyze the volumes.

Action

No action is required.

HWN024203W The volume *storage_volume_name* cannot be converted or moved because the target pools do not have sufficient available space or the target pool types are incorrect for the operation.

Explanation

The volumes cannot be converted or moved because there is not enough space in the target pools or the target pool types are not allowed for the operation.

Action

Add more capacity to the target pools or choose target pools with sufficient available space.

HWN025000I Storage pool *value* in storage system *value* has storage from different types of back-end storage systems. Back-end disk data cannot be determined.

Explanation

Since the back-end storage is from different types of storage systems, back-end disk data like the disk type, raid type and number of disks cannot be determined for the storage pool.

Action

Back-end storage details for the storage pool can be set by editing the data in the Back-end Storage tab of the storage pool properties panel in the IBM Spectrum Control Web-based GUI.

HWN025001I Storage pool *value* in storage system *value* has storage from unknown back-end storage system(s). Back-end disk data cannot be determined.

Explanation

Since the back-end storage is from storage system(s) unknown to IBM Spectrum Control, back-end disk data like the disk type, raid type and number of disks cannot be determined for the storage pool.

Action

Back-end storage details for the storage pool can be set by editing the data in the Back-end Storage tab of the storage pool properties panel in the IBM Spectrum Control Web-based GUI.

HWN025002I Storage pool value in storage system value has storage from multiple back-end storage systems or from multiple pools in a single storage system. Back-end disk data cannot be determined.

Explanation

Since the back-end storage is from multiple storage systems or multiple pools in a single storage system, back-end disk data like the disk type, raid type and number of disks cannot be determined for the storage pool.

Action

Back-end storage details for the storage pool can be set by editing the data in the Back-end Storage tab of the storage pool properties panel in the IBM Spectrum Control Web-based GUI.

HWN025003I Storage pool value in storage system value has storage from a back-end storage pool with multiple disk types. Back-end disk data cannot be determined.

Explanation

Since the back-end storage is from a storage pool with multiple disk types, back-end disk data like the raid type cannot be determined for the storage pool.

Action

Back-end storage details for the storage pool can be set by editing the data in the Back-end Storage tab of the storage pool properties panel in the IBM Spectrum Control Web-based GUI.

HWN025004I Storage pool value in storage system value has storage from a back-end storage pool with a mixed raid type. Back-end disk data cannot be determined.

Explanation

Since the back-end storage is from a storage pool with a mixed raid type, back-end disk data like the raid type cannot be determined for the storage pool.

Action

Back-end storage details for the storage pool can be set by editing the data in the Back-end Storage tab of the storage pool properties panel in the IBM Spectrum Control Web-based GUI.

HWN025005I Storage pool value in storage system value has storage from a back-end storage pool with multiple raid types. Back-end disk data cannot be determined.

Explanation

Since the back-end storage is from a storage pool with multiple raid types, back-end disk data like the raid type cannot be determined for the storage pool.

Action

Back-end storage details for the storage pool can be set by editing the data in the Back-end Storage tab of the storage pool properties panel in the IBM Spectrum Control Web-based GUI.

HWN025006I Storage pool value in storage system value has storage from back-end disks of unknown type. Back-end disk data cannot be determined.

Explanation

Since the back-end storage is from disks of unknown type, back-end disk data cannot be determined for the storage pool.

Action

Back-end storage details for the storage pool can be set by editing the data in the Back-end Storage tab of the storage pool properties panel in the IBM Spectrum Control Web-based GUI.

HWN025007I Storage pool value in storage system value has storage from unknown number of back-end disks. Back-end disk data cannot be determined.

Explanation

Since the back-end storage is from unknown number of disks, back-end disk data cannot be determined for the storage pool.

Action

Back-end storage details for the storage pool can be set by editing the data in the Back-end Storage tab of the storage pool properties panel in the IBM Spectrum Control Web-based GUI.

HWN025008I Storage pool value in storage system value has storage from back-end disks with unknown raid type. Back-end disk data cannot be determined.

Explanation

Since the back-end storage is from disks with unknown raid type, back-end disk data cannot be determined for the storage pool.

Action

Back-end storage details for the storage pool can be set by editing the data in the Back-end Storage tab of the storage pool properties panel in the IBM Spectrum Control Web-based GUI.

HWN025009E Connection to Data Server failed. Make sure Data Server is up.

Explanation

Connection to Data Server cannot be established.

Action

If Data Server is down, start the Data Server. If it is up, make sure you can connect to the Data Server through other operations. If the problem persists, please contact IBM Software Support.

Related reference

-  [Getting support](#)

HWN025011W All of the target ports for the storage system are used for the provisioning request. The request might take a long amount of time.

Explanation

Target port information was not supplied and the Device Server attempted to limit provisioning to the target ports that are on the same fabric as the host. The Device Server cannot determine which target ports are on the same fabric with the host.

Action

The provisioning tasks can complete faster if the fabrics are managed by IBM Spectrum Control. Add the fabrics to IBM Spectrum Control. If the fabrics are already managed by IBM Spectrum Control, check that data collection for the fabrics completed successfully.

HWN025010I Collecting parent pool volumes for storage system: *storage system identification.*

Explanation

The probe is collecting data to gather the relationships of the parent pool volumes to the thin-provisioned volumes.

Action

No action is required.

HWN025011E The port *the target port ID* has a usage restriction which prevents it from being used as a target port for volume assignment.

Explanation

The target port has a usage restriction that does not allow it to be used for volume assignment to initiators. The usage restriction may be in place if the storage system is a virtualizer and the storage system port is dedicated to be used as an initiator for backend storage.

Action

Specify storage system ports that do not have usage restrictions.

HWN025012E The invocation of CIM method *ExposePaths* failed on SMI-S provider *Name of SMI-S provider* . The return code is *Return code of method*.

Explanation

The extrinsic CIM method that was invoked on the given SMI-S provider failed. The standard return codes are as follows:

- 1 - Not Supported
- 2 - Unspecified Error
- 3 - Timeout
- 4 - Failed
- 5 - Invalid Parameter
- 4097 - Invalid logical unit ID
- 4098 - Invalid initiator port ID
- 4099 - Invalid target port ID
- 4100 - Invalid permission
- 4101 - Target/initiator combination already exposed
- 4102 - Requested logical unit number in use
- 4103 - Maximum Map Count Exceeded

Action

Check the details of the request to ensure that there is no problematic input. Otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN025013E The invocation of CIM method HidePaths failed on SMI-S provider *Name of SMI-S provider* . The return code is *Return code of method*.

Explanation

The extrinsic CIM method that was invoked on the given SMI-S provider failed. The standard return codes are as follows:

- 1 - Not Supported
- 2 - Unspecified Error
- 3 - Timeout
- 4 - Failed
- 5 - Invalid Parameter
- 4097 - Invalid logical unit ID
- 4098 - Invalid initiator port ID
- 4099 - Invalid target port ID
- 4100 - Target/initiator combination not exposed

Action

Check the details of the request to ensure that there is no problematic input. Otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN025014E The invocation of CIM method CreateOrModifyElementFromStoragePool failed on SMI-S provider *Name of SMI-S provider* . The return code is *Return code of method*.

Explanation

The extrinsic CIM method that was invoked on the given SMI-S provider failed. The standard return codes are as follows:

- 1 - Not Supported
- 2 - Unknown
- 3 - Timeout
- 4 - Failed
- 5 - Invalid Parameter
- 6 - In Use
- 4097 - Size Not Supported

Action

Check the details of the request to ensure that there is no problematic input. Otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN025015E The invocation of CIM method ReturnToStoragePool failed on SMI-S provider *Name of SMI-S provider* . The return code is *Return code of method*.

Explanation

The extrinsic CIM method that was invoked on the given SMI-S provider failed. The standard return codes are as follows:

- 1 - Not Supported
- 2 - Unknown
- 3 - Timeout
- 4 - Failed
- 5 - Invalid Parameter
- 6 - In Use

Action

Return code 5 can also indicate that the volume is already removed. Check the details of the request to ensure that there is no problematic input. Otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN025016E The invocation of CIM method DeleteStorageHardwareID failed on SMI-S provider *Name of SMI-S provider* . The return code is *Return code of method*.

Explanation

The extrinsic CIM method that was invoked on the given SMI-S provider failed. The standard return codes are as follows:

- 1 - Not Supported
- 2 - Unspecified Error
- 3 - Timeout
- 4 - Failed
- 5 - Invalid Parameter
- 4096 - Specified instance not found

Action

Check the details of the request to ensure that there is no problematic input. Otherwise contact IBM support for more assistance.

Related reference

- [Getting support](#)

HWN025017E A CLI command failed. Check the logs from *EP working dir*.

Explanation

A command issued to CLI failed.

Action

Check the logs from the specified location.

HWN025018E An error occurred when attempting to parse the file *File name*.

Explanation

An error occurred when attempting to parse the specified file.

Action

Verify the correct file was specified or the file is in the correct format.

HWN025019E The requested operation failed. Check the logs from *EP working dir*.

Explanation

The requested operation failed.

Action

Check the logs from the specified location.

HWN025020E The volume cannot be created. The volume of size *Size* in pool *Pool* on storage system *Subsystem* cannot be created. The pool might already have the maximum number of volumes allowed.

Explanation

The volume cannot be created. This pool might have already reached the maximum number of volumes allowed.

Action

Change the provisioning request to make use of pools that have not reached the maximum number of volumes allowed.

HWN025021E Unable to resolve the address for the device because the request was not processed by the data collector.

Explanation

The data collector did not respond to the server in the allotted time. The data collector might not be running or it might not be able to connect to the server.

Action

Verify that the data collector is running and that it can connect to the server.

HWN025022E The data collection detected storage system *New Subsystem* with serial number *new serial number* instead of expected serial number *expected serial number*.

Explanation

The data source reported data for a different storage system as expected, because the a new serial number of the device was returned.

This can happen when a storage system is replaced by a new one with same ip address and credentials.

The new storage system is added and the storage system that was managed previously is shown as undetectable.

Action

Please verify status and the data source information of the storage systems and update them accordingly.

For the new storage system create data collection and performance data collection scheules as needed.

HWN025025I Starting the task to send the report for schedule *Schedule Id* by email.

Explanation

The task to send the the report by email at the frequency configured in the report is being created.

HWN025026I The *report title* report is being created.

Explanation

Based on the settings that were configured, the report is being created.

HWN025027I The *report title* report with ID *report id* is being sent by email to the reports recipients.

Explanation

The report that was created is being sent to the recipients.

HWN025028I The *report title* report with ID *report id* was sent by email to the reports recipients.

Explanation

The task that was created to send the report by email was completed successfully.

HWN025029E Can't retrieve the configured settings of the report for schedule *Schedule Id* .

Explanation

The settings that were used to configure the report were deleted.

Action

To generate the report, a new report must be created and configured.

HWN025030E The report can't be sent because the email server was not configured.

Explanation

The email server was not set up or was not set up correctly.

Action

Set up the email server for sending reports and test the connection.

Related reference

- [Getting support](#)

HWN025031E Can't send the *report title* report with ID *report id* by email because of the following error: *reported_error*.

Explanation

The task failed because an error was generated.

Action

Check the logs for more details about the errors.

Related reference

- [Getting support](#)

HWN025031I To view the report, choose HTML as the message format or use an email application that supports HTML message formats.

Explanation

The report contained in the e-mail can only displayed if the e-mail application supports HTML and has HTML support enabled.

Action

Enable HTML support in your the e-mail client you are using to view this report.

HWN025032E Job failed during post processing of collected data from the data source.

Explanation

The probe returned as failure because of a fatal error during post processing.

Action

Check the logs for more details about the error.

HWN025033E Failed to send the *report name* report for schedule *Schedule Id*.

Explanation

Failed to send the report by e-mail.

HWN025034I Created *number_of_servers* agentless servers automatically.

Explanation

The probe of the storage system discovered one or more host connections for which no corresponding agentless server was defined. The agentless servers for those host connections were created automatically during probe post-processing.

Action

No action is required.

HWN025035I Removed *number_of_servers* agentless servers automatically.

Explanation

The probe of the storage system discovered that one or more agentless servers were no longer associated with any storage system host connections. These agentless servers are obsolete and were removed automatically during probe post-processing.

Action

No action is required.

HWN025036E Can't save the report in the directory.

Explanation

Write access is required to save reports in the directory.

Action

Check that the access permissions for the directory are set correctly .

HWN025037E Can't save the report because the path specifies a file name instead of a directory name.

Explanation

The path must specify the name of the directory that is to be used to save the reports.

Action

Enter the path to a valid directory.

HWN025038E Can't save the report, because the directory doesn't exist.

Explanation

The report can't be saved in the directory that was specified.

Action

Check that the correct path to the directory was specified and that the access permissions for all of the directories in the path are set correctly.

HWN025039E Can't save the report because the directory doesn't have enough disk space.

Explanation

The directory that was specified doesn't have enough available space for saving the report.

Action

Allocate more space to the directory or specify a directory with sufficient space to save the reports in.

HWN025040I *The report title report with ID report id is being saved as report file name in the full path directory.*

Explanation

The report is being saved in the directory that was specified by the user.

HWN025041I *The report title report with ID report id was saved as report file name in the full path directory.*

Explanation

The report was saved in the directory that was specified by the user.

HWN099990I *The method name of the Device Server method of the device server returned return value @(execution context information) .*

Explanation

This is a device server informational trace message, intended for IBM development and support purposes.

Action

No action is required.

HWN099991I *info trace message@(execution context information)*

Explanation

This is a device server informational trace message, intended for IBM development and support purposes.

Action

No action is required.

HWN099992W *warning trace message@(execution context information)*

Explanation

This is a device server warning trace message, intended for IBM development and support purposes.

Action

No action is required.

HWN099993E *error/exception trace message @(execution context information)*

Explanation

This is a device server error/exception trace message, intended for IBM development and support purposes.

Action

No action is required.

HWN099994I *An object of class name of the class has been instantiated @(execution context information).*

Explanation

This is a device server informational trace message, intended for IBM development and support purposes. It indicates the creation of a Java object.

Action

No action is required.

HWN099995I *|=== class name.method name entry, parameter(s): parameter value(s) @(execution context information).*

Explanation

This is a device server informational trace message, intended for IBM development and support purposes. It reports a Java method entry.

Action

No action is required.

HWN099996I *===| class name.method name exit, return value: method return value (execution time in milliseconds) @(execution context information).*

Explanation

This is a device server informational trace message, intended for IBM development and support purposes. It reports a Java method completion.

Action

No action is required.

HWN099997I *External service name of the (DM) external service will be invoked with parameter(s) parameter value(s)@(execution context information).*

Explanation

This is a device server informational trace message, intended for IBM development and support purposes. It reports invocation of an service outside of DM, for example a CIMClient call.

Action

No action is required.

HWN099998I *Invocation of external service name of the (DM) external service returned result invocation result@(execution context information).*

Explanation

This is a device server informational trace message, intended for IBM development and support purposes. It reports the result of the invocation of an service outside of DM, for example a CIMClient call.

Action

No action is required.

HWN099999I The method *name of the device server method* of the device server was invoked with parameters *invocation parameters@(execution context information)*.

Explanation

This is a device server informational trace message, intended for IBM development and support purposes.

Action

No action is required.

HWN200000I Probe of switch *switch_name* completed successfully.

Explanation

The probe of the switch completed successfully.

Action

No action is required.

HWN200001I Started post-processing tasks after data was collected for switch *switch_name*.

Explanation

Data for the switch was collected by using REST, and the probe is complete. Started post-processing tasks for the probe.

Action

No action is required.

HWN6001I Operation *operation* completed successfully.

Explanation

The operation completed successfully.

Action

No response is necessary. This message is informational only.

HWN6002I Unable to set up NLS message file processing.

Explanation

The system was unable to set up message file processing for other languages.

Action

No response is necessary. This message is informational only.

HWN6003E Unable to set up tracing.

Explanation

The agent was unable to set up tracing.

Action

Try the same operation again. If the problem persists, contact IBM service.

HWN6004E Operation operation failed.

Explanation

The operation specified failed.

Action

Try the same operation again. If the problem persists, enable high level tracing as explained in the Installation and Configuration Guide and contact IBM service.

HWN6005E Unknown operation operation.

Explanation

An unknown operation was attempted.

Action

Try the same operation again. If the problem persists, enable high level tracing as explained in the Installation and Configuration Guide and contact IBM service.

HWN6006E Could not initialize connection, rc is rc

Explanation

A connection could not be initialized. Return code from the subsystem was rc.

Action

Try the same operation again. Verify connectivity with the subsystem. If the problem persists, enable high level tracing as explained in the Installation and Configuration Guide and contact IBM service.

HWN6007E Could not parse command arguments: arg

Explanation

The command arguments were not parsed successfully.

Action

Try the same operation again. If the problem persists, enable high level tracing as explained in the Installation and Configuration Guide and contact IBM service.

HWN6008E Error processing command: command

Explanation

There was an error processing the command.

Action

Try the same operation again. If the problem persists, enable high level tracing as explained in the Installation and Configuration Guide and contact IBM service.

HWN6009E Missing 'operation' property in input file

Explanation

The input specified was missing an input property.

Action

Specify the missing property and try the same operation again. If the problem persists, enable high level tracing as explained in the Installation and Configuration Guide and contact IBM service.

HWN6010I Task *arg* completed successfully

Explanation

The specified task completed successfully.

Action

No response is necessary. This message is informational only.

HWN6011E Task *arg* failed

Explanation

The specified task failed.

Action

Try the same operation again. If the problem persists, enable high level tracing as explained in the Installation and Configuration Guide and contact IBM service.

HWN6012E Cannot connect to this IP, switching to *IP*

Explanation

The current IP did not allow connection. The task will switch to the one specified.

Action

Try the same operation again. If the problem persists, enable high level tracing as explained in the Installation and Configuration Guide and contact IBM service.

HWN6013E An IBM XIV CLI command failed. The error is *arg*.

Explanation

The command issued to the XIV CLI failed with the error specified.

Action

Try the same operation again. If the problem persists, look up the XIV error. Contact IBM service if the issue cannot be resolved.

HWN6014I Command *arg* completed successfully

Explanation

A method was called with wrong attributes.

Action

No response is necessary. This message is informational only.

HWN6015E Command *command* failed.

Explanation

The specified command failed.

Action

Try the same operation again. If the problem persists, enable high level tracing as explained in the Installation and Configuration Guide and contact IBM service.

HWN6016I Connected with IP address *IP*

Explanation

Connected with the XIV at the specified IP address

Action

No response is necessary. This message is informational only.

HWN6017I Started creation of volume with size *size* in pool *pool*.

Explanation

A method was called with wrong attributes.

Action

No response is necessary. This message is informational only.

HWN6018I Volume creation completed successfully. New volume *volume* created with size *size* in pool *pool*.

Explanation

A method was called with wrong attributes.

Action

No response is necessary. This message is informational only.

HWN6019I Started deletion of volume *volume* in pool *pool*.

Explanation

Deletion of the specified volume started.

Action

No response is necessary. This message is informational only.

HWN6020I Volume deletion completed successfully. Volume *wolume* deleted in pool *pool*

Explanation

Volume deletion completed successfully. The specified volume was deleted from the specified pool.

Action

No response is necessary. This message is informational only.

HWN6021I Started creation of host *host* with initiator ports *ports*

Explanation

Started host creation.

Action

No response is necessary. This message is informational only.

HWN6022I Finished creation of host *host* with initiator ports *ports*

Explanation

Finished host creation.

Action

No response is necessary. This message is informational only.

HWN6023I Started assignment of volume *volume* to host *host*.

Explanation

Started assignment of the volume to the host.

Action

No response is necessary. This message is informational only.

HWN6024I Finished assignment of volume *volume* to host *host*.

Explanation

Completed assignment of volume to host.

Action

No response is necessary. This message is informational only.

HWN6025I Started unassignment of volume *volume* from host *host*.

Explanation

Started removing volume from host.

Action

No response is necessary. This message is informational only.

HWN6026I Finished unassignment of volume *volume* from host *host*

Explanation

Volume removal completed.

Action

No response is necessary. This message is informational only.

HWNEP0001I Successfully persisted *number of count* instances.

Explanation

A Native API action has been executed successfully. No error condition has been encountered.

Action

No action is required.

HWNEP0002E The probe failed as the data collector couldn't write to its output file, *value*.

Explanation

Information couldn't be collected about the resource as the output file for the data collector was missing or couldn't be accessed.

Action

Try the following actions::

1. Verify that the file system is not full.
2. Verify that the output file exists and can be accessed. The default location for the output file is:
 - Windows: TPC_installation_directory\device\log
 - Linux and AIX: TPC_installation_directory/device/log
3. For more information about the error, check the log files for the device. The default location for the log files is:
 - Windows: TPC_installation_directory\device\log
 - Linux and AIX: TPC_installation_directory/device/log

HWNEP0003E A DS8000 ESSNI command failed. The error code is *error_code*.

Explanation

An ESSNI command failed. No subsequent ESSNI commands were issued, but any commands issued previously were successful. Please check the DS8000 Information Center for more information on the ESSNI error.

Action

Check the DS8000 Information Center for details on the cause and recommended response. Some additional detail may be available in the IBM Spectrum Control message and trace logs. Next, try running the IBM Spectrum Control operation again. Should the problem persist, contact IBM support.

HWNEP0004I Started creation of volume group *volume_group*.

Explanation

Volume Group creation has started. The log will provide information about the subsequent result.

Action

No action is required.

HWNEP0005I Finished creation of volume group *volume_group* with subsystem volume group number *number* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0006I Started adding volumes, with serial numbers *volume_list*, to subsystem volume group *volume_group_number* .

Explanation

Volumes specified will be added to the specified volume group. The log will provide information about the subsequent result.

Action

No action is required.

HWNEP0007I Finished adding volumes to volume group.

Explanation

The task succeeded.

Action

No action is required.

HWNEP0008I Started assignment of host *host* on subsystem *subsystem* to volume group *volume_group*.

Explanation

The host will be associated with the specified volume group, providing it with access to all volumes within the volume group.

Action

No action is required.

HWNEP0009I Finished assigning *host* on subsystem *subsystem* to volume group *volume_group*.

Explanation

The task succeeded.

Action

No action is required.

HWNEP0010I Started removing volumes, with serial numbers *volume_list*, from subsystem volume group *volume_group_number* .

Explanation

Volumes specified will be removed from the specified volume group. The log will provide information about the subsequent result.

Action

No action is required.

HWNEP0011I Finished removing volumes, with serial numbers *volume_list*, from subsystem volume group *volume_group_number* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0012I Increased virtual capacity of storage pool *storage_pool* on subsystem *subsystem* to size *size* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0013I Collecting pools for storage system *storage system identification*.

Explanation

The probe is currently querying the ESSNI server for the extent pools of the storage system.

Action

No action is required.

HWNEP0014I Collecting volumes for *lss* logical subsystems on storage system *storage system identification*.

Explanation

The probe is currently querying the ESSNI server for information regarding logical subsystems and their volumes.

Action

No action is required.

HWNEP0015I Collecting volume groups on storage system *storage system identification.*

Explanation

The probe is currently querying the ESSNI server for the subsystem's volume groups.

Action

No action is required.

HWNEP0016I Collecting hosts on storage system *storage system identification.*

Explanation

The probe is currently querying the ESSNI server for the subsystem's hosts and the volume group to which they are assigned.

Action

No action is required.

HWNEP0017I *value* Hosts Found.

Explanation

This status message is to inform you of the total number of hosts found for this storage system.

Action

No action is required.

HWNEP0018I Launching external process for devices *devices.*

Explanation

An external process has been started for the devices specified. The log will provide information about the subsequent result.

Action

No action is required.

HWNEP0019I External process for devices *devices* completed successfully.

Explanation

An external process has for the devices specified has completed successfully.

Action

No action required.

HWNEP0020E Could not create connection to NAPI *The IP for the NAPI..*

Explanation

No connection to the NAPI could be established.

Action

Check the External Process logs for further details.

HWNEP0021E ESSNI API query for Space Efficient Volume failed with ESSNI code *ESSNI Code*. Data from ESSNI is considered suspect.

Explanation

The ESSNI API was unable to successfully complete the requested query. This indicates a problem within the subsystem and thus the probe will terminate as any data received may be unreliable.

Action

This is usually due to an issue within the NI Services or the LPARs themselves. If the DS gives no indication as to the failure, try restarting one and then the other to see if the problem is resolved.

HWNEP0022I Started deletion of volume group with number *volume_group_number*.

Explanation

Volume group deletion was started. The log will provide information about the subsequent result.

Action

No action is required.

HWNEP0023I Finished deletion of volume group with number *volume_group_number*.

Explanation

The volume group was successfully deleted.

Action

No action is required.

HWNEP0100I Probing Volumes for Storage System: *value*

Explanation

The probe is finding the volumes (VDisks) for this storage system.

Action

No action is required.

HWNEP0101I Number of Volumes currently found: *value*. Continuing to probe Volumes.

Explanation

The probe is finding the volumes (VDisks). This status update is to inform of the number of volumes (VDisks) that have been processed at this point during the probe.

Action

No action is required.

HWNEP0102I *value* Volumes found.

Explanation

This is the total number of Volumes (VDisks) found on the storage system.

Action

No action is required.

HWNEP0103I Probing Configured Disks for Storage System: *value*.

Explanation

The probe is finding the configured disks (MDisks) for this storage system.

Action

No action is required.

HWNEP0104I Number of Configured Disks Found Currently: *value*. Continuing to Probe Disks.

Explanation

The probe is finding the configured disks (MDisks). This status update is to inform how many configured disks (MDisks) have been processed at this point during the probe.

Action

No action is required.

HWNEP0105I *value* Configured Disks Found.

Explanation

This status message to inform you of the total number of configured disks (MDisks) found for this storage system or storage pool.

Action

No action is required.

HWNEP0106I Probing Views of Host Initiator access to Volumes.

Explanation

The probe is finding the Host Initiator access to Volumes.

Action

Check logs for SQLExceptions logged for ServiceUtils.getConnection().

HWNEP0107I Finished probing Views.

Explanation

The probe for finding the Host Initiator access to Volumes is complete.

Action

Check logs for SQLExceptions logged for ServiceUtils.getConnection().

HWNEP0108I Initializing Probe for storage system *storage system identification*.

Explanation

Probe is being initialized.

Action

No Action is required

HWNEP0109I Probe for storage system *storage system identification* completed.

Explanation

Probe is completed.

Action

No action is required.

HWNEP0110I Collecting Nodes and fibre channel ports for storage system *storage system identification*.

Explanation

The probe is currently collection the Nodes and fibre channel ports of the storage system.

Action

No action is required.

HWNEP0111E The connection to the storage device failed. The error code is *error_code*.

Explanation

The External Process connection to the storage device failed with the specified error code. Please check the log details for more info.

Error codes :

- 0 : There is no connection for the specified device
- 2 : No SSH server found on the device
- 3 : Unsupported version
- 4 : The connection to the device failed
- 5 : Authentication failed
- 6 : Unknown host
- 7 : The passphrase is wrong
- 8 : The passphrase is missing
- 9 : Unknown error
- 10: ESSNI not available
- 11: Private key not found
- 12: Invalid format for the private key
- 49: Unable to establish a connection to the device through http port 80

Action

Verify the connection to the storage device and the user credentials. Check the message and trace logs to get to the root cause.

HWNEP0113E The cluster IP address is not specified in the configuration file.

Explanation

The cluster IP address is not specified in the configuration file.

Action

Run the job again. Should the problem persist, contact IBM support.

HWNEP0114E The trustore location is not specified in the configuration file.

Explanation

The trustore location is not specified in the configuration file.

Action

Run the job again. Should the problem persist, contact IBM support.

HWNEP0115E The IBM Spectrum Control data is out of synch with the device configuration and a re-probe is required for device *device name* .

Explanation

The IBM Spectrum Control data is out of synch with the device configuration and a re-probe is required for this device.

Action

Run the probe again. After probe has completed successful run the command again. Ensure that no other configuration changes are performed on the device. If the problem persists, please contact IBM support.

HWNEP0116E The user configured for the subsystem *subsystem name* is not permitted to perform the requested action.

Explanation

The user configured for the current subsystem is not permitted to perform the requested action.

Action

The user that is configured in IBM Spectrum Control to manage the subsystem or is authenticated using the keyfile does not have the proper rights to perform the requested action. This has to be corrected in the subsystem configuration.

HWNEP0117E The virtual disk (VDisk)-to-host mapping was not created because the volume *vdiskName* is already mapped to the *hostName* host for the Device *deviceName*

Explanation

The virtual disk (VDisk)-to-host mapping could not be created because this VDisk is already mapped to this host.

Action

Run the discovery or probe again. Check the message and trace logs to get to the root cause. If the problem persists, contact IBM support.

HWNEP0115I Starting Control Process for storage system *storage system identification*.

Explanation

The specified control process has been started.

Action

No Action is required

HWNEP0116I Started deletion of volume *VolumeID* on subsystem *Subsystem* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0117I Volume deletion completed successfully. Volume *VolumeID* on subsystem *Subsystem* was deleted.

Explanation

The task succeeded.

Action

No action is required.

HWNEP0118I Started adding Managed Disk(s) *Managed Disk ID* to Managed-disk group *Managed Disk group name* on subsystem *Subsystem*.

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0119I Finished adding Managed Disk(s) *Managed Disk ID* to Managed-disk group *Managed Disk group name* on subsystem *Subsystem*.

Explanation

The task succeeded.

Action

No action is required.

HWNEP0120I Started creation of volume with size *Size* in pool *Pool* on subsystem *Subsystem*

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0121I Volume creation completed successfully. New volume *VolumeID* created with size *Size* in pool *Pool* on subsystem *Subsystem* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0122I Started assignment of volume *VolumeID* on subsystem *Subsystem* to initiator port *Initiator Port* on host *Host* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0123I Finished assignment of volume *VolumeID* on subsystem *Subsystem* to initiator port *Initiator Port* on host *Host Name* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0124I Started unassignment of volume *VolumeID* on subsystem *Subsystem* from initiator port *Initiator Port* on host *Host Name* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0125I Finished unassignment of volume *VolumeID* on subsystem *Subsystem* from initiator port *Initiator Port* on host *Host Name* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0126I Started creation of host *host name* on subsystem *Subsystem* with initiator ports *WWPNs* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0127I Finished creation of host *host name* on subsystem *Subsystem* with initiator ports *WWPNs* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0128I Host name *hostName* already exists for the WWPNs *wwpns* on the device *Subsystem*

Explanation

The host name already exists on the device

Action

Run the discovery or probe again. Check the message and trace logs to get to the root cause. If the problem persists, contact IBM support.

HWNEP0129E The operation failed because the device returned unexpected values.

Explanation

Check configuration and connection to the device interface and make sure the device is up and running properly.

Action

An unexpected error occur in the communication with the storage device. Ensure that the device and connection to the device is working properly. Check the message and trace logs to get to the root cause. Run the operation again. If the problem persists, contact IBM support.

HWNEP0130E A IBM XIV CLI command failed. The error is *error_code*.

Explanation

A IBM XIV Command Line Interface command failed. None of the operations that are part of this request were processed successfully. Please check the detailed error messages..

Action

Check the message and trace logs to get to the root cause. Lookup the device specific error code in the device documentation. Run the operation again. Should the problem persist, contact IBM support.

HWNEP0131I The host definition for host *host name* on subsystem *Subsystem* contains additional Hostports *WWPNs* that will also be assigned to Volume *VolumeID* .

Explanation

The selected hostport is part of a host definition on the subsystem that contains additional hostports. These additional hostports will also be assigned to the volume and can access the volume.

Action

If the volume must not be accessible through the additional hostports use the subsystem element manager and change create separate host definitions for the hostports.

HWNEP0132E The unassignment of Volume *VolumeID* from hostport *WWPN* failed because the definition for host *host name* on subsystem *Subsystem* contains additional hostports *WWPNs* .

Explanation

The hostport is part of a host definition on the subsystem that contains additional hostports. Unassigning the additional hostports could cause access loss to the volume .

Action

Either specify all hostports of the host definition for unassignment command or change the host definition on the subsystem using the device's element manager.

HWNEP0133E Error invoking the external process for device *device name* .

Explanation

There was an error invoking the external process.

Action

Check the logs for errors and retry the operation.

HWNEP0134E Following exception occurred: *exception* .

Explanation

There was an unknown error occurred during processing of the function.

Action

Check the logs for errors and retry the operation.

HWNEP0135E External process failed with error code *error code* .

Explanation

The external process failed.

Action

Check the logs for errors and retry the operation.

HWNEP0136E Error connecting to *IP address* with user ID *user ID* .

Explanation

There was an unknown error connecting to a device like subsystem, server, switch, etc..

Action

Make sure the device is up and running.

HWNEP0137I Job *job ID* submitted for device *device name* .

Explanation

A new job was submitted for the device.

Action

None.

HWNEP0138I External process was successfully executed for device *device name* .

Explanation

A job was successfully executed.

Action

None.

HWNEP0139I An instruction was issued to add a copy of the *volume_name* *volume_size*-byte volume in the *pool_name* pool on the *storage_system_name* storage system.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNEP0140I The copy of the *volume_name* *volume_size*-byte volume with the copy ID of *VolumeID* in the *pool_name* pool on the *storage_system_name* storage system was added successfully.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNEP0141I Probing Internal Drives for Storage System: *value*.

Explanation

The probe is finding the internal drives for this storage system.

Action

No action is required.

HWNEP0142I Number of Internal Drives Found Currently: *value*. Continuing to Probe Internal Drives.

Explanation

The probe is finding the internal drives. This status update is to inform how many internal drives have been processed at this point during the probe.

Action

No action is required.

HWNEP0143I *value* Internal Drives Found.

Explanation

This status message to inform you of the total number of internal drives found for this storage system.

Action

No action is required.

HWNEP0144I Probing Pools for Storage System: *value*.

Explanation

The probe is finding the pools (MDisk groups) for this storage system.

Action

No action is required.

HWNEP0145I Number of Pools Found Currently: *value*. Continuing to Probe Pools.

Explanation

The probe is finding the pools (MDisk groups). This status update is to inform how many pools (MDisk groups) have been processed at this point during the probe.

Action

No action is required.

HWNEP0146I *value* Pools Found.

Explanation

This status message to inform you of the total number of pools (MDisk groups) found for this storage system.

Action

No action is required.

HWNEP0147I Collecting asset and status information about *storage_system_id* storage system.

Explanation

The data collection schedule is starting.

Action

No action is required.

HWNEP0148I Collecting cluster information for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for information about clusters.

Action

No action is required.

HWNEP0149I Collecting file system exports for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for information about file system exports.

Action

No action is required.

HWNEP0150I Collecting nodes for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for information about nodes.

Action

No action is required.

HWNEP0151I Collecting file systems for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for information about file systems.

Action

No action is required.

HWNEP0152I Collecting pools for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for information about storage pools.

Action

No action is required.

HWNEP0153I Collecting file system storage for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for information about file system storage.

Action

No action is required.

HWNEP0154I Collecting filesets for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for information about filesets.

Action

No action is required.

HWNEP0155I Collecting links between file systems and nodes for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for links between file systems and nodes.

Action

No action is required.

HWNEP0156I Collecting quotas for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for information about quotas.

Action

No action is required.

HWNEP0157I Collecting file system snapshots for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for information about file system snapshots.

Action

No action is required.

HWNEP0158I Collecting capacity for *file_system_id* file system.

Explanation

The data collection schedule is querying the storage system for information about file system capacity.

Action

No action is required.

HWNEP0159I Creating the export *export name* on cluster *cluster name* .

Explanation

The export is being created. Check the log file for status.

HWNEP0160I The export *export name* on cluster *cluster name* with path *export path* was created.

Explanation

The export was created.

HWNEP0161I The export *export name* on cluster *cluster name* is being changed.

Explanation

The export is being changed. Check the log file for status.

HWNEP0162I The export *export name* on cluster *cluster name* was changed.

Explanation

The export was changed.

HWNEP0163I Setting quota *quota type* - *quota name* on file system *file system name* .

Explanation

The set quota task has started. The log file will contain information on progress.

Action

No action is required.

HWNEP0164I Quota *quota type* - *quota name* on file system *file system name* has been created.

Explanation

The set quota task has completed.

Action

No action is required.

HWNEP0165I Checking quota on file system *file system name* .

Explanation

The check quota task has started. The log file will contain information on progress.

Action

No action is required.

HWNEP0166I Quota on file system *file system name* has been checked.

Explanation

The check quota task has completed.

Action

No action is required.

HWNEP0167I The export *export name* on cluster *cluster name* is being removed.

Explanation

The export is being removed. Check the log file for status.

HWNEP0168I The export *export name* on cluster *cluster name* was removed.

Explanation

The export was removed. Check the log file for status.

HWNEP0169E Command: *command* did not complete. *IBM SONAS CLI message*

Explanation

The IBM SONAS CLI command did not complete.

Action

Verify the IBM SONAS CLI command syntax is valid.

HWNEP0170I Creating fileset *fileset name* on file system *files system name* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0171I Successfully created fileset *fileset name* on file system *file system name* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0172I Removing fileset *fileset name* on file system *files system name* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0173I Successfully removed fileset *fileset name* on file system *file system name* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0174I Modifying fileset *fileset name* on file system *files system name* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0175I Successfully modified fileset *fileset name* on file system *file system name* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0176I Creating file system *file system* on cluster *cluster name* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0177I Successfully created file system *file system* on cluster *cluster name* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0178I Changing file system *file system* on cluster *cluster name* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0179I Successfully changed file system *file system* on cluster *cluster name* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0180I Removing file system *file system* on cluster *cluster name* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0181I Successfully removed file system *file system* on cluster *cluster name* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0182I Mounting file system *file system* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0183I Successfully mounted file system *file system* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0184I Unmounting file system *file system* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0185I Successfully unmounted file system *file system* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0186I Linking fileset *fileset* on file system *file system* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0187I Successfully linked fileset *fileset* on file system *file system* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0188I Unlinking fileset *fileset* on file system *file system* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0189I Successfully unlinked fileset *fileset* on file system *file system* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0190E The IBM Spectrum Control server could not connect to *IP address* using the SSH protocol.

Explanation

The IBM Spectrum Control server could not create a network connection to the specified IBM Storwize V7000 Unified/IBM SONAS device using the SSH protocol.

Action

Verify that the IBM Storwize V7000 Unified/IBM SONAS device is running and that the connection is not being blocked by a firewall.

HWNEP0191E The IBM Spectrum Control server could not authenticate with *IP address* using the SSH protocol.

Explanation

The IBM Spectrum Control server could not create an SSH connection to the specified IBM Storwize V7000 Unified/IBM SONAS device due to an authentication failure.

Action

Verify that the IBM Spectrum Control server is configured with the correct credentials for the IBM Storwize V7000 Unified/IBM SONAS device.

HWNEP0192E The IBM Spectrum Control server could not execute a command on the IBM Storwize V7000 Unified/IBM SONAS device at *IP address* .

Explanation

The IBM Spectrum Control server could not remotely execute a command on the specified IBM Storwize V7000 Unified/IBM SONAS device using the SSH protocol.

Action

Check the log files for additional details.

HWNEP0193E The *command name* command failed because the following command executed on the NAS device failed with the return code *return code* : *command* returned: *command output*

Explanation

The IBM Spectrum Control command failed because the command that was executed on the NAS device failed.

Action

Check the diagnostics section of the manpage of the failed command for additional details.

HWNEP0195I modify fileset

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0196I change export

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0197I create export

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0198I remove export

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0199I create fileset

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0200I link fileset

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0201I remove fileset

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0202I unlink fileset

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0203I change filesystem

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0204I create filesystem

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0205I mount filesystem

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0206I remove filesystem

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0207I unmount filesystem

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0208I check quota

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0209I set quota

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0210I probe

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0211W The *command name* command completed, however during post-processing the following command executed on the NAS device failed with the return code *return code* : *command* returned:

command output As a result, the IBM Spectrum Control database is now out of sync with the current state of the NAS device.

Explanation

The requested command completed, but during post-processing one of the commands executed on the NAS device failed. Post-processing commands are used to update the IBM Spectrum Control database so this failure has caused the IBM Spectrum Control database to be out of sync with the current state of the device.

Action

Probe the device to update the IBM Spectrum Control database and check the diagnostics section of the manpage of the failed command for additional details on why the post-processing command failed.

HWNEP0212I create disk in modifying file system

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0213I Started deletion of host *host name* on subsystem *Subsystem* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0214I Finished deletion of host *host name* on subsystem *Subsystem* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0215I Collecting cache information for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for cache information.

Action

No action is required.

HWNEP0216I remove cached source

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0217I create cached node

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0218I remove cached node

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0219I create cache

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0220I remove cache

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0221I modify cache source

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0222I Creating cache source *cache_source_name* on cluster *file_system_name*.

Explanation

The task was started. Consult the log for details on the task's progress.

Action

No action is required.

HWNEP0223I Created cache source *cache_source_name* on cluster *file_system_name*.

Explanation

The task completed.

Action

No action is required.

HWNEP0224I Removing cache source *cache_source_name* on cluster *file_system_name*.

Explanation

The task was started. Consult the log for details on the task's progress.

Action

No action is required.

HWNEP0225I Removed cache source *cache_source_name* on cluster *file_system_name*.

Explanation

The task completed.

Action

No action is required.

HWNEP0226I Modifying cache source *cache_source_name* on cluster *file_system_name*.

Explanation

The task was started. Consult the log for details on the task's progress.

Action

No action is required.

HWNEP0227I Modified cache *cache_source_name* on cluster *file_system_name*.

Explanation

The task completed.

Action

No action is required.

HWNEP0228I Creating cache *cache_name* on file system *file_system_name*.

Explanation

The task was started. Consult the log for details on the task's progress.

Action

No action is required.

HWNEP0229I Created cache *cache_name* on file system *file_system_name*.

Explanation

The task completed.

Action

No action is required.

HWNEP0230I Removing cache *cache_name* on file system *file_system_name*.

Explanation

The task was started. Consult the log for details on the task's progress.

Action

No action is required.

HWNEP0231I Removed cache *cache_name* on file system *file_system_name*.

Explanation

The task completed.

Action

No action is required.

HWNEP0232I Modifying cache *cache_name* on file system
file_system_name.

Explanation

The task was started. Consult the log for details on the task's progress.

Action

No action is required.

HWNEP0233I Modified cache *cache_name* on file system
file_system_name.

Explanation

The task completed.

Action

No action is required.

HWNEP0234I modify cache

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0235I create cached source

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0236I Configuring nodes *node_names* as cached nodes.

Explanation

The task was started. Consult the log for details on the task's progress.

Action

No action is required.

HWNEP0237I Configured nodes *node_names* as cached nodes.

Explanation

The task completed.

Action

No action is required.

HWNEP0238I Unconfiguring cached nodes *node_names*.

Explanation

The task was started. Consult the log for details on the task's progress.

Action

No action is required.

HWNEP0239I Unconfigured cached nodes *node_names*.

Explanation

The task completed.

Action

No action is required.

HWNEP0240I Executed control operation on cache *cache_name* on filesystem *file_system_name* .

Explanation

The task completed.

Action

No action is required.

HWNEP0241I control cache

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0242I run prepop

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0243I list prepop

Explanation

Name of the command which will be used as an insert to messages HWNEP0193E and HWNEP0211W

Action

None.

HWNEP0244I Retrieving cache prepopulation status for file system *file system name* .

Explanation

The list prepop task has started. The log file will contain information on progress.

Action

No action is required.

HWNEP0245I Cache prepopulation status for file system *file system name* has been retrieved.

Explanation

The check list prepop has completed.

Action

No action is required.

HWNEP0246I Prepopulate cache data for fileset *fileset_name* on file system *file_system_name* using policy *policy_name*.

Explanation

The task was started. Consult the log for details on the task's progress.

Action

No action is required.

HWNEP0247I Command to pre populate cached data for fileset *fileset_name* was successful.

Explanation

The task completed.

Action

No action is required.

HWNEP0248W An error was encountered while parsing protocol options for export *export_name*. The options were not persisted, the probe will continue.

Explanation

Export protocol options are not in the expected format.

Action

Check the message and trace logs to get to the root cause.

HWNEP0249W The connection to the storage device failed. The error code is *error_code*.

Explanation

The External Process connection to the storage device failed with the specified error code. Please check the log details for more info.

Error codes :

- 0 : There is no connection for the specified device
- 2 : No SSH server found on the device
- 3 : Unsupported version
- 4 : The connection to the device failed
- 5 : Authentication failed
- 6 : Unknown host
- 7 : The passphrase is wrong
- 8 : The passphrase is missing
- 9 : Unknown error
- 10: ESSNI not available
- 11: Private key not found
- 12: Invalid format for the private key

Action

Verify the connection to the storage device and the user credentials. Check the message and trace logs to get to the root cause.

HWNEP0250I Started adding initiator port(s) *initiator ports* to host *host name* on subsystem *subsystem* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0251I Finished adding initiator port(s) *initiator ports* to host *host name* on subsystem *subsystem* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0252W A CLI command completed with warning. The warning message is : *warning_message*

Explanation

A Command Line Interface command completed with a warning message. Even if the operation completed, it might be possible to encounter problems when using the results of the operation. Please check the detailed warning message which can be shown in a different language since storage device might not be set to or support the same language as IBM Spectrum Control does.

Action

Check the warning message to see what the problem was.

HWNEP0253W Volume creation completed with warning. New volume *VolumeID* created with size *Size* in pool *Pool* on subsystem *Subsystem* .

Explanation

The operation completed with warning.

Action

Check the log files for details.

HWNEP0254W Volume deletion completed with warning. Volume *VolumeID* on subsystem *Subsystem* was deleted.

Explanation

The operation completed with warning.

Action

Check the log files for details.

HWNEP0255I The task to execute the recommendations for optimizing the volumes on the storage system with an ID of *storage_system_id* was paused.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNEP0256I The task for optimizing the volumes on the storage system with an ID of *storage_system_id* was canceled.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNEP0257I The task for optimizing the volumes on the storage system with an ID of *storage_system_id* was resumed.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNEP0258E The optimization task cannot be paused because the synchronization rate for the volume cannot be reset. The ID of the volume is *volume_id* and the ID of the storage system is *storage_system_id*.

Explanation

The command that was issued to the storage virtualizer to set the synchronization rate of the volume did not complete.

Action

See the log file in the device/log directory for more information about the problem.

HWNEP0259E The optimization task cannot be resumed because the synchronization rate for the volume cannot be reset. The ID of the volume is *volume_id* and the ID of the storage system is *storage_system_id*.

Explanation

The command that was issued to the storage virtualizer to set the synchronization rate of the volume did not complete.

Action

See the log file in the device/log directory for more information about the problem.

HWNEP0260I Started creation of host port *host port name* on storage system *Storage System* with initiator port *WWPN* .

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0261I Finished creation of host port *host port name* on storage system *Storage System* with initiator port *WWPN* .

Explanation

The task succeeded.

Action

No action is required.

HWNEP0262E The recommendation for the *volume_name* volume was not implemented because the command that was issued by the storage virtualizer returned the following error: *error_message*

Explanation

To determine the cause of the error, read the error message that was generated by the storage virtualizer. If the storage virtualizer does not support the same language that IBM Spectrum Control supports, the error message might be shown in a different language.

Action

Rerun the task to implement the recommendations. If an error is still generated, complete the following actions:

- To learn more about the error, check the documentation for the storage resource .
- To resolve the issue, complete the recommended actions.
- To generate new recommendations, probe the storage resource and rerun the wizard to optimize the volumes.

If you cannot resolve the issue, contact IBM Support.

HWNEP0263I The synchronization of the *volume_name* volume with the volume copy was successful.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNEP0264E The synchronization of the *volume_name* volume with the volume copy was unsuccessful.

Explanation

The instruction to synchronize the volume with the volume copy failed.

Action

Check the log messages and the status of the storage resource to resolve the issues that you identify. Manual cleanup might be required to restore the volume.

HWNEP0265E The CLI command that was issued for the *storage_system_name* storage system failed and generated the following error: *error_message*

Explanation

To determine the cause of the error, read the error message that was generated by the storage system. If the storage system does not support the same language that IBM Spectrum Control supports, the error message might be shown in a different language.

Action

Rerun the task to implement the recommendations. If an error is still generated, complete the following actions:

- To learn more about the error, check the documentation for the storage resource .
- To resolve the issue, complete the recommended actions.
- To generate new recommendations, probe the storage resource and rerun the wizard to optimize the volumes.

If you cannot resolve the issue, contact IBM Support.

HWNEP0266I Started expanding the capacity of volume *volume* on subsystem *subsystem* from *oldsize* to *newsize* bytes.

Explanation

The task was started. The log will contain information about the task progress.

Action

No action is required.

HWNEP0267I Finished expanding the capacity of volume *volume* on subsystem *subsystem* to *newsize* bytes.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNEP0268E The server operating system or version is not supported by IBM Spectrum Control for IBM Spectrum Scale.

Explanation

The server operating systems and versions that are supported by IBM Spectrum Control for IBM Spectrum Scale are listed in the interoperability matrix document.

Action

See the IBM Spectrum Control interoperability matrix at the following URL: <http://www.ibm.com/support/docview.wss?uid=swg21386446>. Go to the Storage section and search for IBM Spectrum Scale.

HWNEP0269E The IBM Spectrum Scale cluster information cannot be displayed. All the nodes in the cluster are down or cannot be contacted.

Explanation

An operation was attempted on a remote cluster node. However, none of the nodes in the cluster are reachable or IBM Spectrum Scale is not accepting commands on any of the nodes.

Action

Ensure that the cluster nodes are available and that the storage product user has sufficient authority to complete the required commands on the nodes. Try the operation again.

HWNEP0270E The switch cannot respond to SNMP queries because of an authentication error.

Explanation

The IBM Spectrum Control server could not create a connection to the switch because of an authentication failure.

Action

Make sure authentication info provided was correct and retry. If problem persists, check the log files for details.

HWNEP0271E The following password decryption exception occurred: *exception*

Explanation

The IBM Spectrum Control server could not decrypt the password.

Action

Resolve the issue that is indicated by the exception.

HWNEP0272E The switch cannot respond to SNMP queries because of the following exception: *exception*

Explanation

This exception might occur for any of the following reasons:

- SNMP is not enabled in the device.
- The given SNMP community string is incorrect.
- Access restrictions.

Action

Resolve the issue that is indicated by the exception.

HWNEP0273E The following exception occurred because the OID format is incorrect: *exception*

Explanation

The OID string was sent in the wrong format.

Action

Resolve the issue that is indicated by the exception.

HWNEP0274E The switch cannot respond to SNMP queries because of a timeout problem.

Explanation

A timeout can occur as a result of throttling enhancements or incorrect authentication parameters.

Action

Ensure that the following conditions are true:

- SNMP is enabled on the switch./
- For SNMPv1, the community string is correct.
- For SNMPv3:
 - The username/password authentication and private protocols are correct.
 - The user is not set on the switch by using one of unsupported encryption types: Triple DES, AES 192 or AES 256.

- There are no network access restrictions.

HWNEP0270I Retrieved the file module address *file_module_address*.

Explanation

The file module address was retrieved from the device.

Action

No action is required.

HWNEP0271I No quota data was collected. Quota limits are not activated for the file systems that are associated with the IBM Spectrum Scale cluster.

Explanation

Information about inode and space usage is collected only for file systems on which quota limits are activated.

Action

No action is required.

HWNEP0272I Collecting file systems that are mounted on the nodes of storage system *storage_system_id*.

Explanation

The data collection schedule is querying the IBM Spectrum Scale cluster for information about file systems that are mounted at the nodes of the cluster.

Action

No action is required.

HWNEP0275W One or more operations failed for the CLI command that was issued for the storage system. The following error was generated: *errorMsg* .

Explanation

The CLI command completed but some errors occurred.

Action

Review the error message that was generated by the storage system. To learn more about the cause of the error and the recommended actions, search the IBM Knowledge Center using the storage system name and version. Check the message and trace logs for more information.

HWNEP0276E Command execution failed because *sudo* is not installed.

Explanation

For users other than root, 'sudo' must be installed and configured to allow command execution on behalf of root.

Action

Install and configure 'sudo' for the user that connects to the system.

HWNEP0277I Commands are executed through 'sudo'.

Explanation

A user other than root is used, so commands are executed through 'sudo' with root authority.

Action

No action is required.

HWNEP0278E User can not execute command through sudo.

Explanation

For users other than root, 'sudo' must be configured to allow command execution on behalf of root.

Action

Add the user to the sudoers list and allow all necessary commands to be executed through sudo.

HWNEP0279I Collecting remote file systems for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for information about remote file systems.

Action

No action is required.

HWNEP0280I Collecting remote file systems that are mounted on the nodes of storage system *storage_system_id*.

Explanation

The data collection schedule is querying the IBM Spectrum Scale cluster for information about remote file systems that are mounted at the nodes of the cluster.

Action

No action is required.

HWNEP0281E The switch is returning corrupted data.

Explanation

The server could not use data from the switch.

Action

Review recent configuration changes and try to run the probe again.

HWNEP0282E Zoning data cannot be collected because there is a transaction in progress on the switch

Explanation

The server could not use data from the switch.

Action

None.

HWNEP0283E VSAN *vsan_name* was not found.

Explanation

The VSAN specified as a parameter can not be found in the switch.

Action

Retry the operation with other VSAN name.

HWNEP0284E No zoning data collected from the switch.

Explanation

Zoning data not retrieved from the switch.

Action

Review recent configuration changes and try to run the probe again.

HWNEP0285E Cannot authenticate to the object storage using the specified user credentials.

Explanation

The IBM Spectrum Control server failed to authenticate to the Keystone identity service using the specified user credentials.

Action

Check the user credentials. If the user credentials for the object storage are different from the user credentials for the GPFS cluster, you must enter both sets of user credentials. Try the operation again.

HWNEP0286E An object storage request failed on the GPFS cluster.

Explanation

The IBM Spectrum Control server cannot connect to the object storage on the GPFS cluster. This error might occur because the object service is disabled or stopped or because the user credentials are invalid.

Action

Verify that the object service is configured correctly and is enabled and started. For more information about configuring the object service for IBM Spectrum Scale, go to the IBM Knowledge Center (<http://www.ibm.com/support/knowledgecenter>).

Verify the user credentials and privileges.

Try the operation again. If the problem persists, contact IBM Software Support.

HWNEP0287E Error when collecting Accounts information from Object Storage Service using REST protocol.

Explanation

The probe failed for an IBM Spectrum Scale storage system that is configured for object storage. This issue might occur if the IBM Spectrum Control Server cannot connect to the OpenStack Swift and Keystone endpoints that are used to access object services.

Action

Complete the following steps to verify that the OpenStack Keystone and Swift services are running and correctly configured:

1. To list the URLs for the Keystone and Swift services, run the following commands as user root on one of the GPFS cluster nodes that are configured for object storage. `~/openrc openstack endpoint list`
2. Ensure that the IBM Spectrum Control server can connect to the IP addresses and host names that are included in the URLs for the Keystone and Swift services. For example, the URL for the Keystone service might be `http://gpfs420proto1:5000/v3`. In this case, the IBM Spectrum Control server must be able to connect to the `gpfs420proto1` host name.

If the problem persists, check the log files for the Device server for error messages that might help determine the problem. The default location of log files for the Device server is:

- Windows: `TPC_installation_directory\device\log`
- Linux and AIX: `TPC_installation_directory/device/log`

HWNEP0288E Error when collecting Containers information from Object Storage Service using REST protocol.

Explanation

Failed when attempting to retrieve Swift Containers from Object Storage service.

Action

Review error messages from External Probe log files of Device Server. Verify OpenStack Swift services are running and correctly configured.

HWNEP0281I Collecting object storage accounts for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the GPFS cluster for information about object storage accounts.

Action

No action is required.

HWNEP0282I Collecting object storage containers for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the GPFS cluster for information about object storage containers.

Action

No action is required.

HWNEP0289E Failed to retrieve container information because the number of containers now exceeds the maximum number of containers that can currently be collected for an account (*MAX Containers*).

Explanation

To prevent probes of GPFS clusters that are configured with object storage from taking an excessive amount of time, the number of containers which can be collected for each account during a probe is limited. This limit can be configured by adjusting the Probe.ObjectContainerMaxPages property.

Action

Use the "tpctool getdscfg" command to retrieve the current value for the Probe.ObjectContainerMaxPages property. Archive or delete any unused containers so that the number of containers for the account no longer exceeds the time limit. Try the operation again.

Alternatively, use the "tpctool setdscfg" command to increase the value for the Probe.ObjectContainerMaxPages property. Then, restart the Device server service and try the operation again.

For more information about using the tpctool command, go to the IBM Knowledge Center (<http://www.ibm.com/support/knowledgecenter>).

HWNEP0290E The probe failed to retrieve object storage account information from the storage system *storage_system_id* because the *userid* user does not have the required authority.

Explanation

The user that is configured to monitor the object storage system does not have authority to retrieve information about accounts.

Action

Check the credentials for the user that is configured to monitor the object storage system. In Keystone, the OpenStack identity service, ensure that the user is assigned the required role and has the authority to retrieve object storage account information. Try the operation again.

HWNEP0291E The probe failed to retrieve object storage container information from the storage system *storage_system_id* because the *userid* user does not have the required authority.

Explanation

The user that is configured to monitor the object storage system does not have authority to retrieve information about containers.

Action

Check the credentials for the user that is configured to monitor the object storage system. In Keystone, the OpenStack identity service, ensure that the user is assigned the required role and has the authority to retrieve object storage container information. Try the operation again.

HWNEP0292E Cannot query the object service for information about accounts and containers as the specified user does not have admin privileges.

Explanation

To query account and container information from the object service, the user must be assigned the admin role in Keystone, the OpenStack identity service. To monitor all accounts and containers, the user must also be assigned the role that is defined in the reseller_admin_role configuration option in the Swift proxy server. The default value for the reseller_admin_role option is ResellerAdmin.

Action

Check the credentials for the user. In Keystone, the OpenStack identity service, ensure that the user is assigned the required role and has the authority to retrieve object storage account and container information. Try the operation again.

HWNEP0293W The probe did not collect information about all the object accounts for the storage system *storage_system_id* as the *userid* user does not have sufficient authority on the storage system.

Explanation

The role assigned to the user that is configured to monitor the storage system does not have sufficient privileges to collect information about all accounts and their associated containers. To monitor all accounts on the object storage system, the user must be configured with the ResellerAdmin role.

Action

Check the role that is assigned to the user that is used to monitor the storage system. To monitor all accounts, assign the ResellerAdmin role to the user in Keystone, the OpenStack Identity service. Try the operation again.

HWNEP0294W An authentication error prevented the switch from responding to SNMP queries regarding the ability of the switch to perform zone control.

Explanation

The IBM Spectrum Control server could not create a connection to the switch when checking the ability to perform zone control because of an authentication failure.

Zone control during provisioning Tasks might not be possible until this is resolved.

Action

Make sure that the switch authentication information is correct and then try the action again. If problem persists, contact IBM Software Support.

HWNEP0295W A timeout prevented the switch from responding to SNMP queries regarding the ability of the switch to perform zone control.

Explanation

A timeout can occur as a result of throttling enhancements or incorrect authentication parameters.

Zone control during provisioning Tasks might not be possible until this is resolved.

Action

Using the switch management application, ensure that the following conditions are true:

- SNMP is enabled on the switch./
- For SNMPv1, the read and write community strings are correct.
- For SNMPv3:
 - The username/password authentication and private protocols are correct.
 - The user is not set on the switch by using one of unsupported encryption types: Triple DES, AES 192 or AES 256.
- There are no network access restrictions.

Try the action again.

HWNEP0296W The switch cannot respond to SNMP queries to check the ability of the switch to perform zone control because of the following exception: *exception*

Explanation

This exception might occur for any of the following reasons:

- SNMP is not enabled in the device.
- The SNMP read or write community strings are incorrect.
- Access restrictions.

Action

Use the switch management application to resolve the issue that is indicated by the exception and try the action again.

HWNEP0297W The switch cannot respond to SNMP queries to check the ability of the switch to perform zone control because of the following exception: *exception*

Explanation

This exception might occur for any of the following reasons:

- SNMP is not enabled in the device.
- The SNMP read or write community strings are incorrect.
- Access restrictions have been defined for the switch.

Action

Use the switch management application to resolve the issue that is indicated by the exception and try the action again.

HWNEP0298I Collecting IBM Cloud Object Storage configuration.

Explanation

The data collection schedule is querying IBM Cloud Object Storage for information about system configuration.

Action

No action is required.

HWNEP0299I Collecting IBM Cloud Object Storage vaults.

Explanation

The data collection schedule is querying IBM Cloud Object Storage for information about vaults.

Action

No action is required.

HWNEP0300I Collecting detailed IBM Cloud Object Storage status.

Explanation

The data collection schedule is querying IBM Cloud Object Storage for information about the health and status of its components.

Action

No action is required.

HWNEP0301W The IP address *ip_address* for the FlashSystem storage system is not the management IP address.

Explanation

The IP address for the FlashSystem storage system that you added is not the management IP address. If you do not use the management IP address, the storage system might be managed incorrectly.

Action

To change the current IP address to the management IP address, right-click the FlashSystem storage system on the Block Storage Systems page and select Connections > Modify Connection.

HWNEP0302I Collecting Transparent Cloud Tiering information for *storage_system_id* storage system.

Explanation

The data collection schedule is querying the storage system for information about its Transparent Cloud Tiering configuration.

Action

No action is required.

HWNEP0303I No Transparent Cloud Tiering configuration was detected on the IBM Spectrum Scale cluster.

Explanation

This problem might occur when Transparent Cloud Tiering is not configured on the IBM Spectrum Scale cluster, or the mmcloudgateway command is not installed on node that is used by IBM Spectrum Control to connect to the cluster.

Action

Verify that Transparent Cloud Tiering is configured on the IBM Spectrum Scale cluster. If Transparent Cloud Tiering is configured but is not detected by IBM Spectrum Control, ensure that the mmcloudgateway command is installed on the cluster node that is used to connect to IBM Spectrum Scale.

HWNEP0304E Cannot connect to IBM Cloud Object Storage.

Explanation

This error might occur because the object service is disabled or stopped or because the user credentials for logging in to IBM Cloud Object Storage are invalid.

Action

To resolve the issue, try the following actions:

- Verify that the object service is configured correctly and is enabled and started.
- Verify that the user credentials and privileges are valid.

Try the action again. If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNEP0305I Collecting disk controllers for storage system *storage system identification*.

Explanation

The probe is currently collecting data for the disk controllers of the storage system.

Action

No action is required.

HWNEP0306I Collecting disks for storage system *storage system identification*.

Explanation

The probe is currently collecting data for the disks of the storage system.

Action

No action is required.

HWNEP0307I Collecting CIFS shares for storage system *storage system identification*.

Explanation

The probe is currently collecting data for the CIFS shares of the storage system.

Action

No action is required.

HWNEP0308I Collecting NFS exports for storage system *storage system identification*.

Explanation

The probe is currently collecting data for the NFS exports of the storage system.

Action

No action is required.

HWNEP0309I The data is being collected by the data collector: *data collector host*.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNEP0310I Discovery found *number* storage systems.

Explanation

The discovery task completed and found the specified number of storage systems that are managed by the SMI-S provider.

Action

No action is required.

HWNEP0311I Probing nodes or directors for *storage system name* storage system.

Explanation

The probe is querying the SMI-S provider for information about nodes or directors on the storage system.

Action

No action is required.

HWNEP0312I Probe found *number* nodes or directors.

Explanation

The probe found the specified number of nodes or directors on the storage system.

Action

No action is required.

HWNEP0313I Probing pools for *storage system name* storage system.

Explanation

The probe is querying the SMI-S provider for information about pools on the storage system.

Action

No action is required.

HWNEP0314I Probe found *number* pools.

Explanation

The probe found the specified number of pools on the storage system.

Action

No action is required.

HWNEP0315I Probing disk groups for *storage system name* storage system.

Explanation

The probe is querying the SMI-S provider for information about disk groups on the storage system.

Action

No action is required.

HWNEP0316I Probe found *number* disk groups.

Explanation

The probe found the specified number of disk groups on the storage system.

Action

No action is required.

HWNEP0317I Probing disks for *storage system name* storage system.

Explanation

The probe is querying the SMI-S provider for information about disks on the storage system.

Action

No action is required.

HWNEP0318I Probe found *number* disks.

Explanation

The probe found the specified number of disks on the storage system.

Action

No action is required.

HWNEP0319I Probing host connections for *storage system name* storage system.

Explanation

The probe is querying the SMI-S provider for information about host connections on the storage system.

Action

No action is required.

HWNEP0320I Probing ports for *storage system name* storage system.

Explanation

The probe is querying the SMI-S provider for information about the fibre channel and ethernet ports on the storage system.

Action

No action is required.

HWNEP0321I Probing volumes for *storage system name* storage system.

Explanation

The probe is querying the SMI-S provider for information about volumes on the storage system.

Action

No action is required.

HWNEP0322I Probe found *number* volumes. Continuing to probe volumes.

Explanation

The probe is querying the SMI-S provider for information about volumes on the storage system. So far, the probe found the specified number of volumes.

Action

No action is required.

HWNEP0323I Probe found *number* volumes for *storage system name* storage system.

Explanation

The probe found the specified number of volumes on the storage system.

Action

No action is required.

HWNEP0324I Probing NAS nodes for *storage system name* storage system.

Explanation

The probe is querying the SMI-S provider for information about NAS nodes on the storage system.

Action

No action is required.

HWNEP0325I Probe found *number* NAS nodes.

Explanation

The probe found the specified number of NAS nodes on the storage system.

Action

No action is required.

HWNEP0326I Probing file systems that are mounted on the NAS nodes of *storage system name* storage system.

Explanation

The probe is querying the SMI-S provider for information about file systems that are mounted on the NAS nodes of the storage system.

Action

No action is required.

HWNEP0327I Probe found *number* file systems.

Explanation

The probe found the specified number of file systems on all of the NAS nodes on the storage system.

Action

No action is required.

HWNEP0328I Probing file system exports for *storage system name* storage system.

Explanation

The probe is querying the SMI-S provider for information about file system exports on the storage system.

Action

No action is required.

HWNEP0329W *profile name version version number* SMI-S Profile is not supported.

Explanation

The minimum required version for the specified SMI-S profile is not supported by the SMI-S provider used. The job now running might fail to collect some information.

Action

Upgrade the SMI-S provider to one that supports the minimum required version. See the IBM Spectrum Control interoperability matrix at the following URL: <a href=<http://www.ibm.com/support/docview.wss?uid=swg27047049>><http://www.ibm.com/support/docview.wss?uid=swg27047049>. Search for the supported agent or provider for your resource.

HWNEP0330E Unable to find minimum required SMI-S profile to proceed with requested task.

Explanation

The SMI-S provider does not support the minimum required version of at least one of the following SMI-S profiles: Array, NAS Head, or Storage Virtualizer.

Action

Upgrade the SMI-S provider to one that supports the minimum required version. See the IBM Spectrum Control interoperability matrix at the following URL: <a href=<http://www.ibm.com/support/docview.wss?uid=swg27047049>><http://www.ibm.com/support/docview.wss?uid=swg27047049>. Search for the supported agent or provider for your resource.

HWNEP0331I Probing copy pair relationships for *storage system name* storage system.

Explanation

The probe is querying the SMI-S provider for information about copy pair relationships on the storage system.

Action

No action is required.

HWNEP0332I Probe found *number* copy pairs.

Explanation

The probe found the specified number of copy pairs on the storage system.

Action

No action is required.

HWNEP1111E There is no connection for the specified device.

Explanation

The External Process connection to the storage device failed because there is no connection for the specified device.

Action

Verify the connection to the storage device and the user credentials. Check the message and trace logs to get to the root cause.

HWNEP1112E No SSH server found on the device.

Explanation

The External Process connection to the storage device failed because the SSH server was not found on the device

Action

Verify the ssh server on the storage device and the user credentials. Check the message and trace logs to get to the root cause.

HWNEP1113E Unsupported version.

Explanation

The External Process connection to the storage device failed because the SSH version is not supported.

Action

Verify the connection to the storage device and the user credentials. Check the message and trace logs to get to the root cause.

HWNEP1114E The connection to the device failed.

Explanation

The External Process connection to the storage device failed.

Action

Verify the connection to the storage device and the user credentials. Check the message and trace logs to get to the root cause.

HWNEP1115E Authentication failed.

Explanation

The External Process connection to the storage device failed because authentication failed.

Action

Verify the connection to the storage device and the user credentials. Check the message and trace logs to get to the root cause.

HWNEP1116E Unknown host.

Explanation

The External Process connection to the storage device failed due to unknown host.

Action

Verify the connection to the storage device and the user credentials. Check the message and trace logs to get to the root cause.

HWNEP1117E The passphrase is wrong.

Explanation

The External Process connection to the storage device failed because the passphrase used is wrong.

Action

Verify the connection to the storage device and the passphrase used. Check the message and trace logs to get to the root cause.

HWNEP1118E The passphrase is missing.

Explanation

The External Process connection to the storage device failed because the passphrase is missing.

Action

Verify the connection to the storage device and the passphrase used. Check the message and trace logs to get to the root cause.

HWNEP1119E Unknown error.

Explanation

The External Process connection to the storage device failed because an unknown error occurred.

Action

Verify the connection to the storage device and the user credentials. Check the message and trace logs to get to the root cause.

HWNEP1120E ESSNI not available.

Explanation

The External Process connection to the storage device failed because ESSNI not available.

Action

Verify the connection to the storage device and the user credentials. Check the message and trace logs to get to the root cause.

HWNEP1121E Private key not found.

Explanation

The External Process connection to the storage device failed because the private key used was not found.

Action

Verify the connection to the storage device and the private key used. Check the message and trace logs to get to the root cause.

HWNEP1122E Invalid format for the private key.

Explanation

The External Process connection to the storage device failed because the format for the private key is invalid.

Action

Verify the connection to the storage device and the private key used. Check the message and trace logs to get to the root cause.

HWNEP1123E Unable to establish a connection to the device through http port 80.

Explanation

The External Process connection to the storage device failed because we were unable to establish a connection to the device through http port 80.

Action

Verify the connection to the storage device . Check the message and trace logs to get to the root cause.

HWNEP1124I Log collection successfully started for *storage system name* storage system.

Explanation

The collection of support logs has successfully started for the storage system

Action

No action is required.

HWNEP1125E The activity requested is already in progress on *storage system name* storage system.

Explanation

The requested activity is already in progress on this subsystem.

Action

Wait until the currently running activity completes before making this request.

HWNEP1126I The support log activity has started successfully *storage system name* storage system.

Explanation

The requested activity has started successfully.

Action

No action is required.

HWNEP0112E The CLI command that was issued for the storage system failed and generated the following error: *error_message*

Explanation

To determine the cause of the error, read the error message that was generated by the storage system. If the storage system does not support the same language that IBM Spectrum Control supports, the error message might be shown in a different language.

Action

Rerun the task to implement the recommendations. If an error is still generated, complete the following actions:

- To learn more about the error, check the documentation for the storage resource .
- To resolve the issue, complete the recommended actions.
- To generate new recommendations, probe the storage resource and rerun the wizard to optimize the volumes.

If you cannot resolve the issue, contact IBM Support.

HWNEP1127I The probe failed to retrieve encryption information from the storage system *storage_system_id* because the *userid* user does not have the required authority.

Explanation

The user that is configured to monitor the HPE 3PAR storage system does not have authority to retrieve information about encryption.

Action

Check the credentials for the user that is configured to monitor the object storage system. Collection encryption information requires a user that has Super or Service permission. Try the operation again.

HWNEP1128E The process failed because it was unable to find the Export Tool. Expected location was *loc of tool*.

Explanation

The Hitachi Export Tool is required to collect Performance data for Hitachi systems. The tool was not found at the location specified.

Action

Check that the tool exists and confirm that the location is correct in the Data Collector setup.properties file. Try the operation again.

HWNEP1129E The process failed because the *userid* or password provided failed to connect to the Export Tool.

Explanation

Performance data collection for Hitachi systems requires a user name and password to be defined in the Hitachi Device Manager for the Export Tool. Connection with the provided user name or password failed.

Action

Check the credentials for the user that is configured to monitor the block storage system. To connect to the storage system for both probe and performance monitoring, ensure that the following credentials match:

- The user name and password that is defined in Hitachi Device Manager.
- The user name and password that is used for Hitachi Command Suite to connect to the device.

HWNEP1130E The process failed because the Hitachi SVP was busy and did not return data or timed out.

Explanation

The hitachi external process was unable to collect data because the Hitachi SVP is too busy.

Action

The Hitachi SVP seems to be too busy to return performance data. Collection will retry at the next performance interval.

HWNPM5412E Performance statistics collection is not enabled.

Explanation

Performance statistics collection is not enabled on either the SMI-S provider or the system it manages.

Action

Enable performance statistics collection on either the SMI-S provider or the system it manages.

HWNEP1131E The process failed because the Hitachi performance interval is set to something other than 1 or 5 minutes.

Explanation

The performance interval on the Hitachi Device Manager must be set to either 1 or 5 minutes.

Action

Using the Device Manager, set the performance interval to 1 or 5 minutes.

HWNEP1132W Can't collect further system information because the device returned unexpected values.

Explanation

All the DS8k CECs are in maintenance mode and can't collect further system information.

Action

Make sure the DS8k CECs aren't in maintenance mode when running probe on the storage.

NAD0001I Connecting to *hostname* using *protocol* protocol.

Explanation

The server is connecting to the Storage Resource Agent at the specified hostname using the specified communication protocol.

NAD0002W Connection to *hostname* failed using *protocol* protocol: *error*.

Explanation

The server failed to connect to the Storage Resource Agent at the specified hostname using the specified communication protocol.

NAD0003I Connected to *hostname* using *protocol* protocol.

Explanation

The server connected to the Storage Resource Agent at the specified hostname using the specified communication protocol.

NAD0005E Connection to *hostname* failed using *protocol* protocol: *error message*.

Explanation

The server failed to connect to the Storage Resource Agent at the specified hostname using the specified communication protocol.

NAD0006E Exception thrown for method *method name*: *error message*.

Explanation

An error occurred while processing the specified method.

NAD0007I Closing connection to *hostname*.

Explanation

Closing the connection to the Storage Resource Agent at the specified hostname.

NAD0008E Invalid protocol *protocol* passed to *method name*.

Explanation

The method does not support this protocol.

NAD0010E Invalid parameter(s) *parameter name* passed to *method name*.

Explanation

Invalid input to the specified method.

NAD0013I Installing GUID on remote machine: *hostname*.

Explanation

Installing the unique identifier on the specified machine.

NAD0014I GUID successfully installed on remote machine: *hostname*.

Explanation

Unique identifier installed on the specified machine.

NAD0018E Command on remote machine: *host name* failed. Error code = *value* executing command *value*.

Explanation

The installation on the remote machine failed.

NAD0019E Parameter *parameter* passed to *method* is null or 0 length.

Explanation

Invalid input to the specified method.

NAD0055E Failed to connect to remote host *host*.

Explanation

The specified host is down or not reachable via network.

NAD0097I Opening connection to *hostname*.

Explanation

Opening the connection to the Storage Resource Agent at the specified hostname.

NAD0180I Installing re-distributable package on .

Explanation

Updating remote machine with Visual Studio re-distributable dll package.

NAD0181I Install of re-distributable package on succeeded.

Explanation

Update of Visual studio 2008 DLL's on remote machine succeeded.

NAD0182E Failed to install re-distributable package on .

Explanation

Failed to update remote machine with Visual Studio dll's.

NAD0186I Trying to locate package TIVguid using pkginfo ...

Explanation

Trying to locate installation folder of package TIVguid using pkginfo.

NAD0187I Package TIVguid is not installed.

Explanation

Package TIVguid is not installed on the system.

NAD0188I Checking TIVguid default install path : *path* ...

Explanation

Trying to detect existing installation of package TIVguid using the default install path : /opt/tivoli/guid.

NAD0259W Unable to determine Storage Resource Agent version on host . Fabric Discovery will not be invoked.

Explanation

Storage Resource Agent version could not be obtained. Since Fabric Discovery is not supported on some early versions of Storage Resource Agent, it will not be invoked.

Action

Check if there are any connectivity issues with the Storage Resource Agent. Also, check the device server message and trace log for more detail. If the problem continues, contact IBM support.

NAD0145E Cannot get version information from agent on host .

Explanation

Storage Resource Agent version could not be obtained.

Action

Check if there are any connectivity issues with the Storage Resource Agent. If the problem continues, contact IBM support.

Related reference

-  [Getting support](#)

NAD0146E The connection to *remote machine* failed because the Remote Execution and Access component was unable to create a temporary directory on the remote machine. Remove unneeded ~CSRI* directories in the remote machine's temporary directory.

Explanation

The connection to the remote machine failed.

Action

Remove unneeded ~CSRI* directories in the remote machine's temporary directory.

NAD0156E The server *host_address* cannot be reached because the host name or IP address is not recognized.

Explanation

The host cannot be contacted because the host is unreachable.

Action

Check the host name specified for the target host. Check if the host name can resolve to an IP address from a command line by using commands such as nslookup or ping.

NAD0157E The server *host_name* cannot be contacted. The server might be down, unreachable due to network problems, or the SSH credentials might be invalid.

Explanation

The server cannot be contacted. The server might be down or unreachable due to other connection failures or the server cannot be contacted using the credentials that are provided.

Action

Check the host name that is specified for the server. Check whether the host name can resolve to an IP address from a command line by using commands such as nslookup or ping. Check whether the user and password that are specified are valid. If SSH certificates are used for authentication, check whether the certificate and the passphrase are valid.

NAD0260I Agent is active.

Explanation

The check to ensure that communication with the agent has completed with success.

NAD0272W The connection to the Storage Resource Agent on *host name* was not established. Retrying using the IP address.

Explanation

The Storage Resource Agent specified on this computer could not be connected using the fully qualified computer name. Connection is retried using the ip address.

Action

Check if the fully qualified computer name can be resolved through the network using network commands.

NAD0274E An SSH certificate *certificate name* already exist.

Explanation

The SSH certificate could not be copied on IBM Spectrum Control server directory as there is another certificate with this name.

Action

Rename the SSH certificate name and try again.

NAD0275E Failed to connect to remote host *hostname and port*. Failed to establish a secure connection.

Explanation

The SSH connection failed due to an unknown SSL error.

Action

Verify the connection is not blocked by a firewall and retry the connection.

NAD0276E Failed to connect to remote host *hostname and port*. Failed to establish a secure connection because the SSL handshake failed.

Explanation

The SSL connection failed because the IBM Spectrum Control server and the storage resource agent could not negotiate the desired level of security. This could happen if the IBM Spectrum Control server or storage resource agent certificate is not trusted, not valid, or expired.

Action

If you have recently updated the certificate on the IBM Spectrum Control server then make sure you have also replaced the certificate on the storage resource agent. Make sure that the system date on both the IBM Spectrum Control server and storage resource agent machine is within validity date range of both certificates.

NAD0277E Failed to connect to remote host *hostname and port*. Failed to establish a secure connection because of an invalid SSL key.

Explanation

The SSL connection failed because of a bad SSL key. This is normally caused by a misconfiguration of the SSL certificate and private key.

Action

Verify that the certificates and private keys on the IBM Spectrum Control server and storage resource agent have been configured correctly.

NAD0278E Failed to connect to remote host *hostname and port*. Failed to establish a secure connection because the identity of the peer could not be verified.

Explanation

The SSL connection failed because the peer was not able to identify itself. This could happen if no certificate is available or the configured cipher suite does not support authorization.

Action

Verify that the certificates on the IBM Spectrum Control server and storage resource agent have been configured correctly. There is no need to investigate the cipher suite because IBM Spectrum Control does not allow this to be configured.

NAD0279E Failed to connect to remote host *hostname and port*. Failed to establish a secure connection because of an SSL protocol error.

Explanation

The SSL connection failed because of an error in the SSL protocol.

Action

Contact IBM support.

Related reference

- [Getting support](#)

NAD0281E The Storage Resource agent cannot be deployed because of insufficient space or other issues on the target system. The error is: *error message*.

Explanation

This problem might occur when there is insufficient space on the target system, or the target file is read-only or is being used by another application.

Action

Ensure that the target file is not open or in read-only mode, and that enough space is available on the target system. Try the action again.

BTAVM2272W Unsupported virtual disk backing info for disk "*Disk name*" of hypervisor *Hypervisor name*, virtual machine "*VM name*": *Virtual disk type*.

Explanation

Currently Spectrum Control supports only the following types of ESX virtual disks: VirtualDiskFlatVer2BackingInfo, VirtualDiskSparseVer2BackingInfo, and VirtualDiskRawDiskMappingVer1BackingInfo.

Action

Contact IBM support.

BTAVM2273W Unable to find file "*File name*" which is the backing device of the virtual disk "*Disk name*" of hypervisor *Hypervisor name*, virtual machine "*VM name*".

Explanation

The backing info of the virtual disk was not found by the ESX probe. This happens when either the virtual disk was created after the latest datastore scan or VM was probed while its snapshots were being deleted.

Because the virtual disk could not be correlated to its backing device, the capacity of the VM will be reported wrong in Spectrum Control.

Action

Rerun the hypervisor probe.

BTAVM2274W Probe of hypervisor *Name of the Hypervisor* completed with warnings.

Explanation

The probe of the hypervisor completed with warnings.

Action

Check the message log to find the warnings reason.

BTAEC - Event correlator messages

- [BTAEC1020W The Device Server cannot listen for Forwarded SNMP Traps on port {0}. Port {1} will be used instead.](#)

BTAEC1020W The Device Server cannot listen for Forwarded SNMP Traps on port {0}. Port {1} will be used instead.

Explanation

The port specified for Forwarded SNMP Traps was already in use, so a new one has been found.

BTAHM - Host manager messages

- [BTAHM2501E The service name Service failed to start due to condition.](#)
- [BTAHM2520E Agent agent name has been marked inactive.](#)
- [BTAHM2521E The agent returned an invalid name.](#)
- [BTAHM2522E agent name is not a known agent.](#)
- [BTAHM2524E The agent returned an invalid port number.](#)
- [BTAHM2525E Agent agent name cannot be removed because it is active.](#)
- [BTAHM2527E Unexpected error java error.](#)
- [BTAHM2528I Agent host name has been marked active.](#)
- [BTAHM2551I An inactive agent agent name has been removed.](#)
- [BTAHM2580I The service name service started.](#)
- [BTAHM2581I The service name service is shut down.](#)

BTAHM2501E The *service name* Service failed to start due to *condition*.

Explanation

An error condition prevented the service from starting.

Action

Correct the condition that prevented the service from starting.

BTAHM2520E Agent *agent name* has been marked inactive.

Explanation

The specified agent is not visible to the manager.

Action

Check to make sure that the SANAgentHostQuery service is up and running on the Managed Host.

BTAHM2521E The agent returned an invalid name.

Explanation

An agent name is not in the proper format, according to java.net.InetAddress.

Action

Make sure that the managed host name has a valid IP address and host name.

BTAHM2522E *agent name* is not a known agent.

Explanation

The specified agent is not known to the host manager.

Action

Run a Fabric discovery and then retry the operation.

BTAHM2524E The agent returned an invalid port number.

Explanation

The port number was non-numeric, or was not in the range 0-65535.

Action

Provide a valid port number in the range of 0-65535.

BTAHM2525E Agent *agent name* cannot be removed because it is active.

Explanation

The agent has not been removed.

Action

Remove the agent from the IP network and wait until the server detects that the agent is gone.

BTAHM2527E Unexpected error *java error*.

Explanation

An internal programming error caused an error.

Action

Save the error message and error log and contact your service representative for assistance in resolving the error.

BTAHM2528I Agent *host name* has been marked active.

Explanation

The specified agent is visible to the host manager.

BTAHM2551I An inactive agent *agent name* has been removed.

Explanation

The inactive agent has been removed.

BTAHM2580I The *service name* service started.

Explanation

The host manager started.

BTAHM2581I The *service name* service is shut down.

Explanation

The host manager is shut down.

BTAIC - Inband change agent messages

- [BTAIC1200E The InbandChangeAgent cannot contact the EventCorrelator.](#)
- [BTAIC1201E An error occurred while reading the InbandEvents file.](#)
- [BTAIC1202E The InbandChangeAgent thread has been interrupted.](#)
- [BTAIC1203E The InbandChangeAgent failed to execute the Event.exe command.](#)
- [BTAIC1204E The AIX protocol driver must be uninstalled to prevent it from interfering with the EventScanner.](#)
- [BTAIC1205E In-band event notification requires at least maintenance level 2 for AIX 5.2.](#)
- [BTAIC1206E In-band event notification requires at least maintenance level 1 for AIX 5.3.](#)
- [BTAIC1207E The version of AIX that is running on this managed host is not supported.](#)

BTAIC1200E The InbandChangeAgent cannot contact the EventCorrelator.

Explanation

An error condition prevented the InbandChangeAgent from contacting the EventCorrelator. One possible cause is that the EventCorrelator service is not active.

Action

Verify that all IBM Spectrum Control services are running. Restart the EventCorrelator service if necessary.

BTAIC1201E An error occurred while reading the InbandEvents file.

Explanation

IBM Spectrum Control encountered an error while reading the InbandEvents file.

Action

Contact IBM customer support if the problem persists.

Related reference

- [Getting support](#)

BTAIC1202E The InbandChangeAgent thread has been interrupted.

Explanation

An error condition interrupted the InbandChangeAgent thread. This might be caused by the InbandChangeAgent service shutting down.

Action

Verify that all IBM Spectrum Control services are running. Restart the InbandChangeAgent service if necessary.

BTAIC1203E The InbandChangeAgent failed to execute the Event.exe command.

Explanation

An error condition prevented the InbandChangeAgent from executing the Event.exe command.

Action

Verify that all IBM Spectrum Control services are running. Restart any services that are not running.

BTAIC1204E The AIX protocol driver must be uninstalled to prevent it from interfering with the EventScanner.

Explanation

The agent does not use the protocol driver for this version of AIX. The protocol driver can prevent the EventScanner from gathering in-band events on an AIX host.

Action

Please uninstall the protocol driver using the instructions in the Planning and Installation Guide. Restart the agent after uninstalling the protocol driver.

BTAIC1205E In-band event notification requires at least maintenance level 2 for AIX 5.2.

Explanation

In-band event notification functionality requires at least maintenance level 2 for AIX 5.2.

Action

Please apply maintenance level 2 or greater to the AIX host and restart the agent.

BTAIC1206E In-band event notification requires at least maintenance level 1 for AIX 5.3.

Explanation

In band event notification functionality requires at least maintenance level 1 for AIX 5.3.

Action

Please apply maintenance level 1 or greater to the AIX host and restart the agent.

BTAIC1207E The version of AIX that is running on this managed host is not supported.

Explanation

This AIX managed host is running an unsupported level of the operating system. The agent might not function properly.

Action

Please contact IBM customer support.

Related reference

-  [Getting support](#)

BTALG - Logging toolkit messages

- [BTALG0001I Logging Toolkit is ready.](#)
- [BTALG0002I log add <logger_name>](#)
- [BTALG0003I -handler=<handler_name>](#)
- [BTALG0004I log debug {on|off}](#)
- [BTALG0005I log get](#)
- [BTALG0006I -filterkey](#)
- [BTALG0007I -locale](#)
- [BTALG0008I -format](#)
- [BTALG0009I -maxfiles](#)
- [BTALG0010I -maxfilesize](#)
- [BTALG0011I log get <object_name>](#)
- [BTALG0012I -filename](#)
- [BTALG0013I -filterkey](#)
- [BTALG0014I -formatter](#)
- [BTALG0015I -handlers](#)
- [BTALG0016I -locale](#)
- [BTALG0017I -logstate](#)
- [BTALG0018I -maxfiles](#)
- [BTALG0019I -maxfilesize](#)
- [BTALG0020I log help](#)
- [BTALG0021I log list](#)
- [BTALG0022I -formatters](#)
- [BTALG0023I -locales](#)
- [BTALG0024I -loggers](#)
- [BTALG0025I -handlers](#)
- [BTALG0026I log remove <logger_name>](#)
- [BTALG0027I -handler=<handler_name>](#)
- [BTALG0028I log set](#)
- [BTALG0029I -defaults](#)
- [BTALG0030I -filterkey <INFO|ERROR|WARN>](#)
- [BTALG0031I -format {plain_text|pdxml}](#)
- [BTALG0032I -locale {<locale>|default}](#)

- [BTALG0033I -maxfiles <max_files>](#)
- [BTALG0034I -maxfilesize <max_file_size>](#)
- [BTALG0035I log set <object_name>](#)
- [BTALG0036I -filename <file_name>](#)
- [BTALG0037I -filterkey <INFO|ERROR|WARN>](#)
- [BTALG0038I -formatter <formatter_name>](#)
- [BTALG0039I -locale <locale>](#)
- [BTALG0040I -logstate {on|off}](#)
- [BTALG0041I -maxfiles <max_files>](#)
- [BTALG0042I -maxfilesize <max_file_size>](#)
- [BTALG0043I Invalid number of parameters.](#)
- [BTALG0044I Invalid option.](#)
- [BTALG0045I Function not supported for native loggers.](#)
- [BTALG0046I Locale is set to locale.](#)
- [BTALG0047I Logger is state.](#)
- [BTALG0048I on](#)
- [BTALG0049I off](#)
- [BTALG0050I Attached handlers are handlers.](#)
- [BTALG0051I Filterkey is set to filterkey.](#)
- [BTALG0052I Format is set to format.](#)
- [BTALG0053I Formatter is set to formatter.](#)
- [BTALG0054I Filename is set to filename.](#)
- [BTALG0055I Maxfiles is set to maxfiles.](#)
- [BTALG0056I Maxfilesize is set to maxfilesize KB.](#)
- [BTALG0057I Locale was set to locale.](#)
- [BTALG0058I Filterkey was set to filterkey.](#)
- [BTALG0059I Invalid format format.](#)
- [BTALG0060I Format was set to format.](#)
- [BTALG0061I Formatter was set to format.](#)
- [BTALG0062I Filename was set to filename.](#)
- [BTALG0063I Maxfiles was set to maxfiles.](#)
- [BTALG0064I Maxfilesize was set to maxfilesize KB.](#)
- [BTALG0065I Logging defaults have been restored.](#)
- [BTALG0066I Failed to update property.](#)
- [BTALG0067I Logger has been turned state.](#)
- [BTALG0068I Invalid option option.](#)
- [BTALG0069I -handler](#)
- [BTALG0070I Handler added successfully.](#)
- [BTALG0071I Failed to add handler.](#)
- [BTALG0072I Handler removed successfully.](#)
- [BTALG0073I Failed to remove handler.](#)
- [BTALG0074I Invalid command command.](#)
- [BTALG0075I Debug is set to state.](#)
- [BTALG0076I plain_text](#)
- [BTALG0077I pdxml](#)
- [BTALG0078I Failed to get property information.](#)
- [BTALG0079I Displays logging properties.](#)
- [BTALG0080I Provides general information on the Logging Service commands.](#)
- [BTALG0081I Defines logging properties.](#)
- [BTALG0082I Adds a handler to the specified logger.](#)
- [BTALG0083I Enables or disables additional logging commands.](#)
- [BTALG0084I Provides a list of loggers, handlers, or formatters.](#)
- [BTALG0085I Removes a handler object.](#)
- [BTALG0086I log add <logger_name> \[option\]](#)
- [BTALG0087I log debug {on|off}](#)
- [BTALG0088I log get \[option\]](#)
- [BTALG0089I log help \[option\]](#)
- [BTALG0090I log list \[option\]](#)
- [BTALG0091I log remove <logger_name> \[option\]](#)
- [BTALG0092I log set \[option\]](#)
- [BTALG0093I IBM Spectrum Control Logging Toolkit for Fabric](#)
- [BTALG0094I Command Line Interface - Version version Release release Level level minor](#)
- [BTALG0095I LOGGING SERVICE COMMANDS](#)
- [BTALG0097I Command](#)
- [BTALG0098I Description](#)
- [BTALG0099I See](#)
- [BTALG0100I -add](#)
- [BTALG0101I -debug](#)
- [BTALG0102I -get](#)
- [BTALG0103I -help](#)
- [BTALG0104I -list](#)
- [BTALG0105I -remove](#)
- [BTALG0106I -set](#)
- [BTALG0107I OPTION](#)
- [BTALG0108I COMMAND](#)
- [BTALG0109I DESCRIPTION](#)
- [BTALG0110I add](#)
- [BTALG0111I debug](#)
- [BTALG0112I get](#)
- [BTALG0113I list](#)

- [BTALG0114I remove](#)
- [BTALG0115I set](#)
- [BTALG0116I Adds the handler to the specified logger.](#)
- [BTALG0117I log get <object_name> \[option\]](#)
- [BTALG0118I Displays the current types of messages that are logged in the log file.](#)
- [BTALG0119I Displays the maximum number of log files to be created.](#)
- [BTALG0120I Displays the maximum file size of the log before a new log file is created.](#)
- [BTALG0121I Displays the format in which messages are saved in the log file.](#)
- [BTALG0122I Displays the current language locale setting in which messages are displayed in the message log file.](#)
- [BTALG0123I Displays the file name associated with the specified handler.](#)
- [BTALG0124I Displays the formatter that is attached to the specified handler.](#)
- [BTALG0125I Displays the handler that is attached to the specified logger.](#)
- [BTALG0126I Displays if the logger is on or off. You must specify a logger for the object name.](#)
- [BTALG0127I Displays the current types of messages that are logged for the specified logger.](#)
- [BTALG0128I Displays the maximum number of log files to be created for the specified handler.](#)
- [BTALG0129I Displays the maximum file size of log files created by the specified handler.](#)
- [BTALG0130I log set <object_name> \[option\]](#)
- [BTALG0131I Logging configuration corrupted. Restoring default configuration.](#)
- [BTALG0132I Displays a list of loggers.](#)
- [BTALG0133I Displays a list of formatters.](#)
- [BTALG0134I Displays a list of handlers.](#)
- [BTALG0135I Removes the handler from the specified logger.](#)
- [BTALG0136I -handlers=<list of handlers>](#)
- [BTALG0137I Specifies the types of messages that will be logged.](#)
- [BTALG0138I Sets the maximum number of log files to be created.](#)
- [BTALG0139I Sets the maximum file size \(in kilobytes\) of the log before a new log file is created.](#)
- [BTALG0140I Sets the format in which messages are saved in the log file.](#)
- [BTALG0142I Sets the file name where the specified handler will output log messages.](#)
- [BTALG0143I Sets the formatter used by the specified handler.](#)
- [BTALG0145I Turns the log on or off. You must specify a logger for the object name.](#)
- [BTALG0146I Specifies the types of messages that will be logged.](#)
- [BTALG0147I Sets the maximum number of log files to be created.](#)
- [BTALG0148I Sets the maximum file size \(in kilobytes\) of the log before a new log file is created.](#)
- [BTALG0149I Resets the logging properties to their default settings.](#)
- [BTALG0150I help](#)
- [BTALG0151I State](#)
- [BTALG0152I Filter](#)
- [BTALG0153I Handlers](#)
- [BTALG0154I Logger](#)
- [BTALG0155E logger is not a valid logger. Failed to update property.](#)

BTALG0001I Logging Toolkit is ready.

Explanation

NoHelp

BTALG0002I log add <logger_name>

Explanation

This is a command line command.

BTALG0003I -handler=<handler_name>

Explanation

This is a command line command.

BTALG0004I log debug {on|off}

Explanation

This is a command line command.

BTALG0005I log get

Explanation

This is a command line command.

BTALG0006I -filterkey

Explanation

This is a command line command.

BTALG0007I -locale

Explanation

This is a command line command.

BTALG0008I -format

Explanation

This is a command line command.

BTALG0009I -maxfiles

Explanation

This is a command line command.

BTALG0010I -maxfilesize

Explanation

This is a command line command.

BTALG0011I log get <object_name>

Explanation

This is a command line command.

BTALG0012I -filename

Explanation

This is a command line command.

BTALG0013I -filterkey

Explanation

This is a command line command.

BTALG0014I -formatter

Explanation

This is a command line command.

BTALG0015I -handlers

Explanation

This is a command line command.

BTALG0016I -locale

Explanation

This is a command line command.

BTALG0017I -logstate

Explanation

This is a command line command.

BTALG0018I -maxfiles

Explanation

This is a command line command.

BTALG0019I -maxfilesize

Explanation

This is a command line command.

BTALG0020I log help

Explanation

This is a command line command.

BTALG0021I log list

Explanation

This is a command line command.

BTALG0022I -formatters

Explanation

This is a command line command.

BTALG0023I -locales

Explanation

This is a command line command.

BTALG0024I -loggers

Explanation

This is a command line command.

BTALG0025I -handlers

Explanation

This is a command line command.

BTALG0026I log remove <logger_name>

Explanation

This is a command line command.

BTALG0027I -handler=<handler_name>

Explanation

This is a command line command.

BTALG0028I log set

Explanation

This is a command line command.

BTALG0029I -defaults

Explanation

This is a command line command.

BTALG0030I -filterkey <INFO | ERROR | WARN>

Explanation

This is a command line command.

BTALG0031I -format {plain_text|pdxml}

Explanation

This is a command line command.

BTALG0032I -locale {<locale>|default}

Explanation

This is a command line command.

BTALG0033I -maxfiles <max_files>

Explanation

This is a command line command.

BTALG0034I -maxfilesize <max_file_size>

Explanation

This is a command line command.

BTALG0035I log set <object_name>

Explanation

This is a command line command.

BTALG0036I -filename <file_name>

Explanation

This is a command line command.

BTALG0037I -filterkey <INFO|ERROR|WARN>

Explanation

This is a command line command.

BTALG0038I -formatter <formatter_name>

Explanation

This is a command line command.

BTALG0039I -locale <locale>

Explanation

This is a command line command.

BTALG0040I -logstate {on|off}

Explanation

This is a command line command.

BTALG0041I -maxfiles <max_files>

Explanation

This is a command line command.

BTALG0042I -maxfilesize <max_file_size>

Explanation

This is a command line command.

BTALG0043I Invalid number of parameters.

Explanation

NoHelp

BTALG0044I Invalid option.

Explanation

NoHelp

BTALG0045I Function not supported for native loggers.

Explanation

NoHelp

BTALG0046I Locale is set to *locale*.

Explanation

NoHelp

BTALG0047I Logger is *state*.

Explanation

NoHelp

BTALG0048I on

Explanation

This is a command line command.

BTALG0049I off

Explanation

This is a command line command.

BTALG0050I Attached handlers are *handlers*.

Explanation

NoHelp

BTALG0051I Filterkey is set to *filterkey*.

Explanation

NoHelp

BTALG0052I Format is set to *format*.

Explanation

NoHelp

BTALG0053I Formatter is set to *formatter*.

Explanation

NoHelp

BTALG0054I Filename is set to *filename*.

Explanation

NoHelp

BTALG0055I Maxfiles is set to *maxfiles*.

Explanation

NoHelp

BTALG0056I Maxfilesize is set to *maxfilesize* KB.

Explanation

NoHelp

BTALG0057I Locale was set to *locale*.

Explanation

NoHelp

BTALG0058I Filterkey was set to *filterkey*.

Explanation

NoHelp

BTALG0059I Invalid format *format*.

Explanation

NoHelp

BTALG0060I Format was set to *format*.

Explanation

NoHelp

BTALG0061I Formatter was set to *format*.

Explanation

NoHelp

BTALG0062I Filename was set to *filename*.

Explanation

NoHelp

BTALG0063I Maxfiles was set to *maxfiles*.

Explanation

NoHelp

BTALG0064I Maxfilesize was set to *maxfilesize* KB.

Explanation

NoHelp

BTALG0065I Logging defaults have been restored.

Explanation

BTALG0066I Failed to update property.

Explanation

NoHelp

BTALG0067I Logger has been turned *state*.

Explanation

NoHelp

BTALG0068I Invalid option *option*.

Explanation

NoHelp

BTALG0069I -handler

Explanation

This is a command line command.

BTALG0070I Handler added successfully.

Explanation

NoHelp

BTALG0071I Failed to add handler.

Explanation

NoHelp

BTALG0072I Handler removed successfully.

Explanation

NoHelp

BTALG0073I Failed to remove handler.

Explanation

NoHelp

BTALG0074I Invalid command *command*.

Explanation

NoHelp

BTALG0075I Debug is set to *state*.

Explanation

NoHelp

BTALG0076I *plain_text*

Explanation

This is a command line command.

BTALG0077I *pdxml*

Explanation

This is a command line command.

BTALG0078I Failed to get property information.

Explanation

NoHelp

BTALG0079I Displays logging properties.

Explanation

NoHelp

BTALG0080I Provides general information on the Logging Service commands.

Explanation

NoHelp

BTALG0081I Defines logging properties.

Explanation

NoHelp

BTALG0082I Adds a handler to the specified logger.

Explanation

NoHelp

BTALG0083I Enables or disables additional logging commands.

Explanation

NoHelp

BTALG0084I Provides a list of loggers, handlers, or formatters.

Explanation

NoHelp

BTALG0085I Removes a handler object.

Explanation

NoHelp

BTALG0086I log add <logger_name> [option]

Explanation

This is a command line command.

BTALG0087I log debug {on|off}

Explanation

This is a command line command.

BTALG0088I log get [option]

Explanation

This is a command line command.

BTALG0089I log help [option]

Explanation

This is a command line command.

BTALG0090I log list [option]

Explanation

This is a command line command.

BTALG0091I log remove <logger_name> [option]

Explanation

This is a command line command.

BTALG0092I log set [option]

Explanation

This is a command line command.

BTALG0093I IBM Spectrum Control Logging Toolkit for Fabric

Explanation

NoHelp

BTALG0094I Command Line Interface - Version *version* Release *release* Level *level* *minor*

Explanation

NoHelp

BTALG0095I LOGGING SERVICE COMMANDS

Explanation

NoHelp

BTALG0097I Command

Explanation

This is a command line command.

BTALG0098I Description

Explanation

This is a command line command.

BTALG0099I See

Explanation

This is a command line command.

BTALG0100I -add

Explanation

This is a command line command.

BTALG0101I -debug

Explanation

This is a command line command.

BTALG0102I -get

Explanation

This is a command line command.

BTALG0103I -help

Explanation

This is a command line command.

BTALG0104I -list

Explanation

This is a command line command.

BTALG0105I -remove

Explanation

This is a command line command.

BTALG0106I -set

Explanation

This is a command line command.

BTALG0107I OPTION

Explanation

This is a command line command.

BTALG0108I COMMAND

Explanation

This is a command line command.

BTALG0109I DESCRIPTION

Explanation

This is a command line command.

BTALG0110I add

Explanation

This is a command line command.

BTALG0111I debug

Explanation

This is a command line command.

BTALG0112I get

Explanation

This is a command line command.

BTALG0113I list

Explanation

This is a command line command.

BTALG0114I remove

Explanation

This is a command line command.

BTALG0115I set

Explanation

This is a command line command.

BTALG0116I Adds the handler to the specified logger.

Explanation

NoHelp

BTALG0117I log get <object_name> [option]

Explanation

This is a command line command.

BTALG0118I Displays the current types of messages that are logged in the log file.

Explanation

NoHelp

BTALG0119I Displays the maximum number of log files to be created.

Explanation

NoHelp

BTALG0120I Displays the maximum file size of the log before a new log file is created.

Explanation

NoHelp

BTALG0121I Displays the format in which messages are saved in the log file.

Explanation

NoHelp

BTALG0122I Displays the current language locale setting in which messages are displayed in the message log file.

Explanation

NoHelp

BTALG0123I Displays the file name associated with the specified handler.

Explanation

NoHelp

BTALG0124I Displays the formatter that is attached to the specified handler.

Explanation

NoHelp

BTALG0125I Displays the handler that is attached to the specified logger.

Explanation

NoHelp

BTALG0126I Displays if the logger is on or off. You must specify a logger for the object name.

Explanation

NoHelp

BTALG0127I Displays the current types of messages that are logged for the specified logger.

Explanation

NoHelp

BTALG0128I Displays the maximum number of log files to be created for the specified handler.

Explanation

NoHelp

BTALG0129I Displays the maximum file size of log files created by the specified handler.

Explanation

NoHelp

BTALG0130I log set <object_name> [option]

Explanation

This is a command line command.

BTALG0131I Logging configuration corrupted. Restoring default configuration.

Explanation

NoHelp

BTALG0132I Displays a list of loggers.

Explanation

NoHelp

BTALG0133I Displays a list of formatters.

Explanation

NoHelp

BTALG0134I Displays a list of handlers.

Explanation

NoHelp

BTALG0135I Removes the handler from the specified logger.

Explanation

NoHelp

BTALG0136I -handlers=<list_of_handlers>

Explanation

This is a command line command.

BTALG0137I Specifies the types of messages that will be logged.

Explanation

NoHelp

BTALG0138I Sets the maximum number of log files to be created.

Explanation

NoHelp

BTALG0139I Sets the maximum file size (in kilobytes) of the log before a new log file is created.

Explanation

NoHelp

BTALG0140I Sets the format in which messages are saved in the log file.

Explanation

NoHelp

BTALG0142I Sets the file name where the specified handler will output log messages.

Explanation

NoHelp

BTALG0143I Sets the formatter used by the specified handler.

Explanation

NoHelp

BTALG0145I Turns the log on or off. You must specify a logger for the object name.

Explanation

NoHelp

BTALG0146I Specifies the types of messages that will be logged.

Explanation

NoHelp

BTALG0147I Sets the maximum number of log files to be created.

Explanation

NoHelp

BTALG0148I Sets the maximum file size (in kilobytes) of the log before a new log file is created.

Explanation

NoHelp

BTALG0149I Resets the logging properties to their default settings.

Explanation

NoHelp

BTALG0150I help

Explanation

This is a command line command.

BTALG0151I State

Explanation

This is a command line command.

BTALG0152I Filter

Explanation

This is a command line command.

BTALG0153I Handlers

Explanation

This is a command line command.

BTALG0154I Logger

Explanation

This is a command line command.

BTALG0155E *logger* is not a valid logger. Failed to update property.

Explanation

NoHelp

BTAMS - Spectrum Control Messaging Service messages

- [BTAMS0500I IBM Spectrum Control MessagingService started successfully.](#)
- [BTAMS0501I The Messaging Service has shutdown.](#)
- [BTAMS0502I Service service name subscribed to topic topic name.](#)
- [BTAMS0503I Event published to topic topic name.](#)
- [BTAMS0504E Messaging Service failed to get a proxy to the service service name.](#)
- [BTAMS0505E Messaging Service could not invoke the onMessage method on service service name.](#)
- [BTAMS0001W Failed to load the configuration for the database exception handler.](#)
- [BTAMS0002I Data server](#)
- [BTAMS0003I Device server](#)
- [BTAMS0004I Unknown](#)
- [BTAMS0005W Failed to update the database pool monitor handler handler name.](#)
- [BTAMS0006W Failed to initialize the database pool monitor handler handler name.](#)

BTAMS0500I IBM Spectrum Control MessagingService started successfully.

Explanation

The Messaging Service has started successfully.

BTAMS0501I The Messaging Service has shutdown.

Explanation

The Messaging Service has been shutdown.

BTAMS0502I Service *service name* subscribed to topic *topic name*.

Explanation

The specified service has subscribed to events published to the specified topic.

BTAMS0503I Event published to topic *topic name*.

Explanation

An event has been published to the specified topic.

BTAMS0504E Messaging Service failed to get a proxy to the *service name* service.

Explanation

Messaging Service could not get a proxy to the service.

Action

Review the message log to determine why the proxy could not be obtained. If this problem persists, enable ServiceManager tracing to assist in determining why the proxy could not be obtained.

BTAMS0505E Messaging Service could not invoke the onMessage method on service *service name*.

Explanation

An exception occurred when invoking the onMessage method on the specified service. The service might not be running.

Action

Verify that the service is up and running.

BTAMS0001W Failed to load the configuration for the database exception handler.

Explanation

Failed to load an optional configuration file for the database exception handler. The default configuration will be used.

Action

No action is required.

BTAMS0002I Data server

Explanation

Identifier for the Data server.

Action

No action is required.

BTAMS0003I Device server

Explanation

Identifier for the Device server.

Action

No action is required.

BTAMS0004I Unknown

Explanation

Identifier that is used if the server type cannot be determined.

Action

No action is required.

BTAMS0005W Failed to update the database pool monitor handler *handler name.*

Explanation

Failed to update the specified database pool monitor.

Action

No action is required.

BTAMS0006W Failed to initialize the database pool monitor handler *handler name.*

Explanation

Failed to initialize the specified database pool monitor.

Action

No action is required.

BTAQE - Spectrum Control Query Engine messages

- [BTAQE1100E Query Engine Event Generator can not start.](#)
- [BTAQE1101E Unable to open the database.](#)
- [BTAQE1102E Unable to close the database.](#)
- [BTAQE1104E The Query Engine check write authority failed.](#)
- [BTAQE1105E Check for QueryEngine Authentication failed.](#)
- [BTAQE1106E The SANQueryEngine thread has been interrupted.](#)
- [BTAQE1109E An error occurred while attempting to save the IP target to the database.](#)
- [BTAQE1110E An error occurred while attempting to delete an IP target from database.](#)
- [BTAQE1111E An error occurred while querying the IP target information from the database.](#)
- [BTAQE1116E Database errors occurred while performing queries on Tasks.](#)
- [BTAQE1117E Database errors occurred while saving the task.](#)
- [BTAQE1118E Errors occurred while resolving InterconnectElement and Port relationship.](#)

- [BTAQE1119E Errors in Topology XML generator.](#)
- [BTAQE1120E Errors in creating an entity.](#)
- [BTAQE1121E Invalid target host IP address.](#)
- [BTAQE1122E Failed to get the SNMP Service proxy.](#)
- [BTAQE1123E Unable to close the event publisher.](#)
- [BTAQE1124E Unable to compress scanner result due to IOException: exception.](#)
- [BTAQE1125E Unable to uncompress the scanner result due to IOException: exception.](#)
- [BTAQE1126I An unidentified port was removed from the scan data.](#)
- [BTAQE1127E An outband scanner failed to capture the scan data.](#)
- [BTAQE1128E An outband scanner failed to save the scan data for benchmark comparison.](#)
- [BTAQE1129E An outband scanner failed to read the benchmark file benchmark file name saved from the previous scan.](#)
- [BTAQE1130E An outband scanner failed to decrypt the password for target target IP.](#)
- [BTAQE1134I The outband agent target address TargetIP address does not respond to Fibre Channel MIB \(previously called the Fibre Alliance MIB\) queries.](#)
- [BTAQE1135E Unable to get the license state from the license server.](#)
- [BTAQE1136E The Query Engine cannot obtain a valid IP address for the host target.](#)
- [BTAQE1137E The Query Engine could not obtain the information for target host target from the database.](#)
- [BTAQE1138E The Query Engine could not obtain the capability information for target host target from the database.](#)
- [BTAQE1139E The Query Engine could not obtain the information for all known target hosts from the database.](#)
- [BTAQE1140E Error creating an event subscriber.](#)
- [BTAQE1141E The Query Engine could not obtain the scheduled scan information from the database.](#)
- [BTAQE1142E The Query Engine could not obtain the list of active scanners from the database.](#)
- [BTAQE1143E The Query Engine could not obtain the list of inactive scanners from the database.](#)
- [BTAQE1144E An error occurred attempting to run the scanner name scanner on the IBM Spectrum Control managed host target.](#)
- [BTAQE1145E The scanner name scanner running on the IBM Spectrum Control managed host target found no SAN.](#)
- [BTAQE1146E The scanner name scanner running on IBM Spectrum Control managed host target found no host-based adapter \(HBA\).](#)
- [BTAQE1147E The scanner name scanner running on IBM Spectrum Control managed host target found no SCSI host-based adapter \(HBA\).](#)
- [BTAQE1149E A scanner overlap condition has occurred for the scanner name scanner on the IBM Spectrum Control managed host target .](#)
- [BTAQE1150I The outband agent target address TargetIP address does not support topology discovery through SNMP Fibre Channel MIB \(previously called the Fibre Alliance MIB\), or Cisco VSAN MIB queries.](#)
- [BTAQE1151I The outband agent with target address TargetIP address has been added.](#)
- [BTAQE1152I The outband agent with target address TargetIP address has been removed.](#)

BTAQE1100E Query Engine Event Generator can not start.

Explanation

The SAN Query Engine Event Generator was unable to start.

Action

Get the SANQueryEngine service trace information from the trace log, and contact IBM customer support.

BTAQE1101E Unable to open the database.

Explanation

IBM Spectrum Control was unable to open the database. There might be a problem with the database or the network connections to the database server.

Action

Make sure that the database is online and working properly and is connected to the network. Stop and restart the manager, which will allow the database to synchronize its data.

BTAQE1102E Unable to close the database.

Explanation

IBM Spectrum Control was unable to close the database. There might be a problem with the database or the network connections to the database server.

Action

Make sure that the database is online and working properly and is connected to the network. Stop and restart the manager, which will allow the database to synchronize its data.

BTAQE1104E The Query Engine check write authority failed.

Explanation

The IBM Spectrum Control query engine does not have the appropriate authority. There might be a problem with the product license file or with the license server, or the current license might have expired.

Action

Make sure the IBM Spectrum Control license has not expired. Contact IBM customer support to get a new license, if necessary.

Related reference

- [Getting support](#)

BTAQE1105E Check for QueryEngine Authentication failed.

Explanation

During a routine authorization check, IBM Spectrum Control encountered an error. There might be a problem with the license file or the license server, or the current license might have expired.

Action

Make sure the IBM Spectrum Control license has not expired. Contact IBM customer support to get a new license, if necessary.

Related reference

- [Getting support](#)

BTAQE1106E The SANQueryEngine thread has been interrupted.

Explanation

There is a problem with the SANQueryEngine service. This might be caused by the service shutting down.

Action

Make sure the SANQueryEngine service is active. If it is not running, you may need to stop and restart the manager to start the service again.

BTAQE1109E An error occurred while attempting to save the IP target to the database.

Explanation

There might be a problem with the database or with the network connections to the database.

Action

Make sure the database is running and is properly connected to the network. Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAQE1110E An error occurred while attempting to delete an IP target from database.

Explanation

There might be a problem with the database or with the network connections to the database.

Action

Make sure the database is running and is properly connected to the network. Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAQE1111E An error occurred while querying the IP target information from the database.

Explanation

There might be a problem with the database or with the network connections to the database.

Action

Make sure the database is running and is properly connected to the network. Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAQE1116E Database errors occurred while performing queries on Tasks .

Explanation

There might be a problem with the database or the network connections to the database.

Action

Make sure the database is running and is properly connected to the network.

BTAQE117E Database errors occurred while saving the task.

Explanation

There might be a problem with the database or the network connections to the database.

Action

Make sure the database is running and is properly connected to the network.

BTAQE1118E Errors occurred while resolving InterconnectElement and Port relationship.

Explanation

Some of the data returned from the scanner could not be resolved.

Action

Get the SANQueryEngine service trace information from the trace log, and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAQE1119E Errors in Topology XML generator.

Explanation

Some of the data generated by the scanner appears to have errors.

Action

Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAQE1120E Errors in creating an entity.

Explanation

Errors occurred while attempting to create an entity from the results of a scan.

Action

Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

BTAQE1121E Invalid target host IP address.

Explanation

The IP address that was returned from the target host appears to be invalid.

Action

Verify the IP address of the target host. Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAQE1122E Failed to get the SNMP Service proxy.

Explanation

The scanner was unable to get the SNMP Service proxy. This might be caused by network problems.

Action

Make sure the network is running properly. Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAQE1123E Unable to close the event publisher.

Explanation

IBM Spectrum Control was unable to close the event publisher.

Action

Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAQE1124E Unable to compress scanner result due to IOException: *exception*.

Explanation

An IOException occurred in the process of compressing the scanner result.

BTAQE1125E Unable to uncompress the scanner result due to IOException: *exception*.

Explanation

An IOException occurred in the process of uncompressing the scanner result.

BTAQE1126I An unidentified port was removed from the scan data.

Explanation

The outband AdvancedTopologyScanner scan data contained a port entity which did not have a port World Wide Name (WWN), or it had a port WWN of 0000000000000000. The port cannot be identified.

Action

This problem can be caused by a switch port being improperly configured. Check the switch port configurations. For more information on the unidentified port, get the SANQueryEngine service trace information from the trace log.

BTAQE1127E An outband scanner failed to capture the scan data.

Explanation

An error condition prevented the scanner from processing the outband scan data.

Action

If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTAQE1128E An outband scanner failed to save the scan data for benchmark comparison.

Explanation

An outband scanner encountered an error while saving the scan data as a benchmark file. This file is created when the scanner is run for the first time. When the scanner is invoked again, the scan data is compared against the saved benchmark data from the previous scan for optimization processing. Benchmark comparison will not be performed for the next scan due to the encountered error.

Action

Check to see if the file system is full on the managed host system. If the problem continues, contact IBM customer support.

Related reference

-  [Getting support](#)

BTAQE1129E An outband scanner failed to read the benchmark file *benchmark file name* saved from the previous scan.

Explanation

An outband scanner encountered an error while reading the benchmark file. This file is created when the scanner is run for the first time. When the scanner is invoked again, the scan data is compared against the saved benchmark data from the previous scan for optimization processing. Benchmark comparison is not performed for the current scan due to the encountered error.

Action

Check to see if the benchmark file name exists in the specified path on the manager system. If the problem continues, contact IBM customer support.

Related reference

-  [Getting support](#)

BTAQE1130E An outband scanner failed to decrypt the password for target *target IP*.

Explanation

An outband scanner encountered an error decrypting an outband agent password to target address target IP . Either the encrypted password does not exist, or it has been corrupted.

Action

Restart the product. If the problem continues, contact IBM customer support.

Related reference

-  [Getting support](#)

BTAQE1134I The outband agent target address *TargetIP address* does not respond to Fibre Channel MIB (previously called the Fibre Alliance MIB) queries.

Explanation

A basic query to the Fibre Channel MIB at IP address TargetIP address did not respond. Either the device does not support the Fibre Channel MIB, or it is not enabled.

Action

If the problem continues, contact IBM customer support.

BTAQE1135E Unable to get the license state from the license server.

Explanation

Unable to get either the license state or the license state value from the license server.

Action

No operator response

BTAQE1136E The Query Engine cannot obtain a valid IP address for the host *target*.

Explanation

The host might have an invalid IP address.

Action

Verify the IP address of the host. Get the SANQueryEngine service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAQE1137E The Query Engine could not obtain the information for target host *target* from the database.

Explanation

An error occurred while attempting to obtain target host information. This could be caused by a network error, or there could be a problem with the database.

Action

Make sure the network is functioning properly. Check to see if the database is online and functioning properly. Retry the operation. If it continues to fail, contact IBM customer support.

Related reference

- [Getting support](#)

BTAQE1138E The Query Engine could not obtain the capability information for target host *target* from the database.

Explanation

An error occurred while attempting to obtain target host capability information. This could be caused by a network error, or there could be a problem with the database.

Action

Make sure the network is functioning properly. Check to see if the database is online and functioning properly. Retry the operation. If it continues to fail, contact IBM customer support.

Related reference

- [🔗 Getting support](#)

BTAQE1139E The Query Engine could not obtain the information for all known target hosts from the database.

Explanation

An error occurred while attempting to obtain target host information. This could be caused by a network error, or there could be a problem with the database.

Action

Make sure the network is functioning properly. Check to see if the database is online and functioning properly. Retry the operation. If it continues to fail, contact IBM customer support.

Related reference

- [🔗 Getting support](#)

BTAQE1140E Error creating an event subscriber.

Explanation

The Query Engine received an error while trying to create an event subscriber. This might have been caused by a service not starting properly. This will not affect system processing.

Action

If this message continues to appear, contact IBM customer support.

Related reference

- [🔗 Getting support](#)

BTAQE1141E The Query Engine could not obtain the scheduled scan information from the database.

Explanation

An error occurred while attempting to obtain scheduled scan information. This could be caused by a network error, or there could be a problem with the database.

Action

Make sure the network is functioning properly. Check to see if the database is online and functioning properly. Retry the operation. If it continues to fail, contact IBM customer support.

Related reference

- [🔗 Getting support](#)

BTAQE1142E The Query Engine could not obtain the list of active scanners from the database.

Explanation

An error occurred while attempting to obtain the list of active scanners. This could be caused by a network error, or there could be a problem with the database.

Action

Make sure the network is functioning properly. Check to see if the database is online and functioning properly. Retry the operation. If it continues to fail, contact IBM customer support.

Related reference

- [Getting support](#)

BTAQE1143E The Query Engine could not obtain the list of inactive scanners from the database.

Explanation

An error occurred while attempting to obtain the list of inactive scanners. This could be caused by a network error, or there could be a problem with the database.

Action

Make sure the network is functioning properly. Check to see if the database is online and functioning properly. Retry the operation. If it continues to fail, contact IBM customer support.

Related reference

- [Getting support](#)

BTAQE1144E An error occurred attempting to run the *scanner name* scanner on the IBM Spectrum Control managed host *target*.

Explanation

IBM Spectrum Control periodically scans each managed host to identify associated devices.

Action

Check the status of the managed hosts. Check the message log file on the managed host for any error information. Restart any hosts that are not running properly.

BTAQE1145E The *scanner name* scanner running on the IBM Spectrum Control managed host *target* found no SAN.

Explanation

An error occurred attempting to run a scan on one of the IBM Spectrum Control managed hosts. The scanner did not find a Storage Area Network (SAN).

Action

Check that the host-based adapter (HBA) is properly connected to the Storage Area Network.

BTAQE1146E The *scanner name* scanner running on IBM Spectrum Control managed host *target* found no host-based adapter (HBA) .

Explanation

An error occurred while attempting to run a scan on one of the IBM Spectrum Control managed hosts. The scanner could not communicate with the HBA.

Action

Check that an HBA is installed in the managed host and is properly connected to the SAN. Also check that the correct level of device drivers are installed.

BTAQE1147E The *scanner name* scanner running on IBM Spectrum Control managed host *target* found no SCSI host-based adapter (HBA) .

Explanation

An error occurred while attempting to run a scan on one of the IBM Spectrum Control managed hosts. The scanner could not communicate with the HBA.

Action

Check that an HBA is installed in the managed host and is properly connected to the SAN. Also check that the correct level of device drivers are installed.

BTAQE1149E A scanner overlap condition has occurred for the *scanner name* scanner on the IBM Spectrum Control managed host *target* .

Explanation

IBM Spectrum Control periodically scans each managed host to identify associated devices. In this case the managed host has determined that more than one scanner process of the same type attempted to run on the managed host at the same time. This condition is known as scanner overlap. The scanner overlap condition is not allowed on the managed host, and should be corrected.

Action

Check the status of the managed host. Check the message log file on the managed host for any error information. Terminate the scanner processes still running on the managed host. Restart the managed host if needed. If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTAQE1150I The outband agent target address *TargetIP address* does not support topology discovery through SNMP Fibre Channel MIB (previously called the Fibre Alliance MIB), or Cisco VSAN MIB queries.

Explanation

A basic query to either the Fibre Channel Management MIB, or the Cisco VSAN MIB to IP address *TargetIP address* did not respond. Either the device does not support the Fibre Channel Management MIB, or the Cisco VSAN MIB, or one of those MIBs is not enabled.

Action

If the problem continues, contact IBM customer support.

BTAQE1151I The outband agent with target address *TargetIP address* has been added.

Explanation

An agent with IP address *TargetIP address* has been added.

Action

This event occurs as a result of adding an outband agent using the Add Fabric Agent panel.

BTAQE1152I The outband agent with target address *TargetIP address* has been removed.

Explanation

An agent with IP address *TargetIP address* has been removed.

Action

This event occurs as a result of removing an outband agent using the Fabric Agents panel.

BTASA - Spectrum Control SAN scanner agent messages

- [BTASA1400E The SAN Agent Scanner failed to execute the inband scanner scanner with the command: scan command.](#)
- [BTASA1401E The SAN Agent Scanner failed to capture the inband scan data.](#)
- [BTASA1403E The SAN Agent Scanner failed to save the scan data for benchmark comparison.](#)
- [BTASA1404E The SAN Agent Scanner failed to read the benchmark file benchmark file saved from the previous scan.](#)
- [BTASA1405E The SAN Agent Scanner failed to retrieve the global unique identifier from Host Query.](#)
- [BTASA1406I The SAN Agent Scanner Service has initialized successfully.](#)
- [BTASA1407I The Inband scanner scanner has started.](#)
- [BTASA1408I The Inband scanner scanner has ended with return code return code.](#)
- [BTASA1409E A scanner overlap condition has occurred on the IBM Spectrum Control managed host.](#)
- [BTASA1420E The GS-3 Zone Control DLL could not be loaded.](#)

BTASA1400E The SAN Agent Scanner failed to execute the inband scanner *scanner* with the command: *scan command*.

Explanation

An error condition prevented the scanner from executing the inband scan.

Action

Verify that the scanner executable scanner exists in the specified path on the managed host system. If the problem continues, contact IBM customer support.

Related reference

- [🔗 Getting support](#)

BTASA1401E The SAN Agent Scanner failed to capture the inband scan data.

Explanation

An error condition prevented the scanner from processing the inband scan data.

Action

If the problem continues, contact IBM customer support.

Related reference

- [🔗 Getting support](#)

BTASA1403E The SAN Agent Scanner failed to save the scan data for benchmark comparison.

Explanation

The scanner encountered an error while saving the scan data as a benchmark file. This file is created when the scanner is run for the first time. When the scanner is invoked again, the scan data is compared against the saved benchmark data from the previous scan for optimization processing. Benchmark comparison is not performed for the next scan due to the encountered error.

Action

Check to see if the file system is full on the managed host system. If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTASA1404E The SAN Agent Scanner failed to read the benchmark file *benchmark file* saved from the previous scan.

Explanation

The scanner encountered an error while reading the benchmark file. This file is created when the scanner is run for the first time. When the scanner is invoked again, the scan data is compared against the saved benchmark data from the previous scan for optimization processing. Benchmark comparison is not performed for the current scan due to the encountered error.

Action

Check to see if the benchmark file *benchmark file* name exists in the specified path on the managed host system. If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTASA1405E The SAN Agent Scanner failed to retrieve the global unique identifier from Host Query.

Explanation

The scanner encountered an error while retrieving the global unique identifier for the managed host agent from Host Query. This identifier is required for the service to initialize.

Action

Verify that the SAN Agent Host Query service is initialized. Try to restart the managed host agent if the SAN Agent Scanner is the only service that is not initialized. Otherwise, follow the instructions in the Install Guide to uninstall and reinstall the agent. If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTASA1406I The SAN Agent Scanner Service has initialized successfully.

Explanation

The scanner service has initialized successfully.

BTASA1407I The Inband scanner *scanner* has started.

Explanation

The specified inband scanner has started on the managed host.

BTASA1408I The Inband scanner *scanner* has ended with return code *return code*.

Explanation

The specified inband scanner has ended on the managed host.

BTASA1409E A scanner overlap condition has occurred on the IBM Spectrum Control managed host.

Explanation

The managed host has determined that more than one scanner process of the same type attempted to run on the managed host at the same time. This condition is known as scanner overlap. The scanner overlap condition is not allowed on the managed host, and should be corrected.

Action

Check the status of the managed host. Check the message log file on the managed host for any error information. Terminate the scanner processes still running on the managed host. Restart the managed host if needed. If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTASA1420E The GS-3 Zone Control DLL could not be loaded.

Explanation

The GS-3 Zone Control DLL was not loaded properly. The Zone Control functions will not operate properly.

Action

Stop and restart the agent. Check that the system has adequate disk and memory available when the agent is started. Check the agent install log for errors. Uninstall and reinstall the agent if necessary. If the problem continues, contact IBM customer support.

Related reference

- [Getting support](#)

BTASD - Fabric User Interface messages

- [BTASD1922E An error occurred while getting the information from device services.](#)
- [BTASD1923E The agent agent identifier is currently agent state. It must be in order to remove it.](#)
- [BTASD1930E Unable to contact zone agent](#)
- [BTASD1931E Unable to contact zoning agent. Token used for contacting zone agent is invalid.](#)
- [BTASD1932E Agent capable of configuring zoning could not be found on this Fabric](#)
- [BTASD1933E Zoning is already being configured on this Fabric. New zoning can not be done until agent is available again](#)
- [BTASD1934E Unable to delete the selected entities.](#)
- [BTASD1935E The delete failed because an agent is still installed on the selected computer or configured in Data Agent or Inband Fabric Agent list.](#)
- [BTASD1936E Unexpected server response message_status= agent response.](#)
- [BTASD1937E A zone name cannot contain the characters '!' '%' '*' or '!' in its name. Brocade zone names also cannot contain '\\$' nor '-'. The first character for a zone name must be alphanumeric. Enter a new name for this zone.](#)
- [BTASD1938E A zone set name cannot contain the characters '!' '%' '*' or '!' in its name. Brocade zone set names also cannot contain '\\$' nor '-'. The first character for a zone set name must be alphanumeric. Enter a new name for this zone set.](#)
- [BTASD1939E Zoning is already being configured by lock owner on this Fabric since lock time. New zoning can not be done until agent is available again.](#)
- [BTASD1940E Zoning is already being configured by lock owner on this Fabric since lock time. New zoning can not be done until agent is available again. Do you want to release the lock from user {Q}?](#)
- [BTASD1941E The lock for the zone control operations has been reset. Do you want to re-obtain the lock and continue zone operations?](#)
- [BTASD1942E Zoning changes cannot be made at this time. Zoning for this fabric is currently locked by lock owner since date.](#)
- [BTASD1943E An alias with the same name already exists. Enter a new name for this alias.](#)
- [BTASD1944E The Alias name field must be filled in to create an alias.](#)
- [BTASD1945E A alias name cannot contain the characters '!' '%' '*' or '!' in its name. Brocade alias names also cannot contain '\\$' nor '-'. The first character for an alias name must be alphanumeric. Enter a new name for this alias.](#)
- [BTASD1946E An alias name cannot begin with a number. Enter a new name for this alias.](#)
- [BTASD1947E An alias, zone and/or zone set in the same configuration can not have the same name. Enter a new name.](#)
- [BTASD1948E This alias does not contain any members. Add a member to this alias.](#)
- [BTASD1949E There are too many members selected. The maximum number of members allowed for this alias is max members. Member\(s\) must be removed before you can continue.](#)
- [BTASD1950E The zone configuration has the maximum number of aliases allowed. The maximum number of aliases is max aliases. An existing alias must be deleted before a new one can be created.](#)
- [BTASD1951E Fabric fabric name has zone count zones with non-standard members: zone names. Zone and ZoneSet changes cannot be applied to zones with nonstandard members.](#)
- [BTASD1952E Zone set zoneset name is active or in activation or deactivation pending status. Deletion of the zone set is not supported for this fabric. After committing the deactivation, this zone set can be deleted in the next zone control session.](#)

- [BTASD1953E Renaming of the active zone set is not supported for this fabric. Deactivate zone set zoneset name before attempting to rename it.](#)
- [BTASD1954E Zoning cannot be done because there is no connection to the SMI-S provider. Reason: reason](#)
- [BTASD2001W Zoning has changed on this fabric since the configuration panel was opened. You may need to run the Fabric discovery/probe again to get the zone information within IBM Spectrum Control synchronized with the fabric. Do you still want to make zoning changes?](#)
- [BTASD2002W This alias is not assigned to any zones. This could result in an error when the zone configuration is applied later. Do you want to continue?](#)
- [BTASD2003W No exclusive fabric-wide lock is available on the switches for fabric fabric name. Other users might be modifying the zoning configuration from outside of IBM Spectrum Control during your zoning operations.](#)
- [BTASD2004W This Out Of Band Agent is already defined with the same parameters. Would you like to save it anyway?](#)
- [BTASD3001I A probe job for fabric fabric name has been submitted. The inactive Zone Definition for this fabric will have the old Zone Definition until the probe job is complete. Wait a few minutes before working with Zone Definition for this fabric.](#)

BTASD1922E An error occurred while getting the information from device services.

Explanation

The device service could not be contacted by the data service. Possible causes are device server is not running. Communication between the data server machine and the device server machine was not possible.

Action

Check to see if the device server is currently operational. Check to see if communication is possible between the data and device server. If it continues to fail, restart the device server then restart the data server.

BTASD1923E The agent *agent identifier* is currently *agent state*. It must be in order to remove it.

Explanation

Agent must be in inactive state in order to remove it.

Action

Stop the inband agent. Try to remove again.

BTASD1930E Unable to contact zone agent

Explanation

Zoning agent could not be found to conduct zoning on this fabric

Action

Determine whether Fabric has a zoneable agent.

Check the status of the device server.

Restart the device server.

BTASD1931E Unable to contact zoning agent. Token used for contacting zone agent is invalid.

Explanation

A token is used to ensure that only one agent can configure zoning at a time. This error can occur when the zoning agent is unable to contact the switch and the token times out.

Action

Try doing zoning operation again. Another agent may have completed its task allowing this one to gain a token.

Determine whether Fabric has a zoneable agent.

Check the status of the device server.

Restart the device server.

BTASD1932E Agent capable of configuring zoning could not be found on this Fabric

Explanation

Zoning a fabric requires an agent capable of zoning be available on the fabric. Typically this is an SRA or an inband fabric agent or a Brocade SMI Agent. A token is used to ensure that only one agent can configure zoning at a time. This error can occur when the zoning agent is unable to contact the switch and the token times out.

Action

Determine if fabric has zoning agent. SRA, Inband fabric, Brocade SMI Agent are typically required.

Check status of fabric agent. If inactive, activate it

Check the status of the device server.

Restart the device server.

BTASD1933E Zoning is already being configured on this Fabric. New zoning can not be done until agent is available again

Explanation

To avoid colliding zoning configurations, two agents can not zone the same fabric at the same time. This error occurs when two separate interfaces attempt to conduct zoning at the same time. The interfaces can be the IBM Spectrum Control GUI, the native switch UI or the IBM Spectrum Control CLI.

Action

Determine if another UI is attempting zoning. Wait for it to complete and try again.

Check the status of the device server.

Restart the device server.

BTASD1934E Unable to delete the selected entities.

Explanation

Action

Check status of fabric agent. If inactive, activate it

Check the status of the device server.

Restart the device server.

BTASD1935E The delete failed because an agent is still installed on the selected computer or configured in Data Agent or Inband Fabric Agent list.

Explanation

Action

Uninstall any Data Server or Device Server agents.

Check the status of the device server.

Restart the device server.

BTASD1936E Unexpected server response message_status= agent response.

Explanation

IBM Spectrum Control has received an unexpected response from the Agent

Action

Run the Check function on the Agent. If the response is successful then retry command.

Otherwise check the IBM Spectrum Control and Agent logs for Agent errors.

BTASD1937E A zone name cannot contain the characters '.',',',',','%','*' or '!' in its name. Brocade zone names also cannot contain '\$' nor '-'. The first character for a zone name must be alphanumeric. Enter a new name for this zone.

Explanation

One or more of the characters in the zone name string is not allowed by the Capabilities for the switch type you are trying to do zone configuration for.

Action

Remove the characters from the zone name that are not allowed for the switch type you are trying to configure.

BTASD1938E A zone set name cannot contain the characters '.',',',',','%','*' or '!' in its name. Brocade zone set names also cannot contain '\$' nor '-'. The first character for a zone set name must be alphanumeric. Enter a new name for this zone set.

Explanation

One or more of the characters in the zone set name string is not allowed for the switch type you are trying to do zone configuration for.

Action

Remove the characters from the zone set name that are not allowed for the switch type you are trying to configure.

BTASD1939E Zoning is already being configured by *lock owner* on this Fabric since *lock time*. New zoning can not be done until agent is available again.

Explanation

Another user has locked the fabric for zoning.

Action

Wait until the Fabric is available for zoning.

BTASD1940E Zoning is already being configured by *lock owner* on this Fabric since *lock time*. New zoning can not be done until agent is available again. Do you want to release the lock from user {0}?

Explanation

Another user has locked the fabric for zoning.

Action

Click YES to reset the lock on the Fabric and to continue.

BTASD1941E The lock for the zone control operations has been reset. Do you want to re-obtain the lock and continue zone operations?

Explanation

The lock on this fabric has been lost.

Action

Choose yes to re-obtain the lock for the Fabric and continue with the zoning operations. If the problem occurs frequently, contact IBM Support.

Related reference

-  [Getting support](#)

BTASD1942E Zoning changes cannot be made at this time. Zoning for this fabric is currently locked by *lock owner* since *date*.

Explanation

Another user is making zone operations and has acquired a lock for this Fabric preventing any other users from making changes at the same time.

Action

Contact the user that has acquired the lock.

BTASD1943E An alias with the same name already exists. Enter a new name for this alias.

Explanation

Fabrics can not have aliases with duplicate names. An alias with the name you entered already exists for this fabric.

Action

Enter a different name for this alias.

BTASD1944E The Alias name field must be filled in to create an alias.

Explanation

The name field in the alias wizard is empty. Every alias must have a name to be created.

Action

Enter a name for this alias.

BTASD1945E A alias name cannot contain the characters '.',',',',',%',',*', or '!' in its name. Brocade alias names also cannot contain '\$' nor '-'. The first character for an alias name must be alphanumeric. Enter a new name for this alias.

Explanation

One or more of the characters in the alias name string is not allowed by the Capabilities for the switch type you are trying to do zone configuration for.

Action

Remove the characters from the alias name that are not allowed for the switch type you are trying to configure. See the switch documentation for details.

BTASD1946E An alias name cannot begin with a number. Enter a new name for this alias.

Explanation

Alias names cannot start with a number.

Action

Change the name to not start with a number.

BTASD1947E An alias, zone and/or zone set in the same configuration can not have the same name. Enter a new name.

Explanation

The alias name entered matches the name for a zone or zone alias in this configuration. This fabric does not support that capability.

Action

Enter a unique name for the alias.

BTASD1948E This alias does not contain any members. Add a member to this alias.

Explanation

This fabric does not support empty aliases. The alias must have members assigned to it.

Action

Select members to be added to the alias and move them to the selected list.

BTASD1949E There are too many members selected. The maximum number of members allowed for this alias is *max members*. Member(s) must be removed before you can continue.

Explanation

Zone Aliases have a maximum number of members allowed. The current alias has reached that limitation.

Action

Remove members from the selected list until the number of members is equal to or less than the maximum number given in the error message.

BTASD1950E The zone configuration has the maximum number of aliases allowed. The maximum number of aliases is *max aliases*. An existing alias must be deleted before a new one can be created.

Explanation

This zone configuration has a maximum number of aliases allowed and has reached that limitation.

Action

Remove aliases from the aliases table so there are less than the maximum number supplied in the error message.

BTASD1951E Fabric *fabric name* has zone count zones with non-standard members: *zone names*. Zone and ZoneSet changes cannot be applied to zones with nonstandard members.

Explanation

This fabric has non-standard zones so IBM Spectrum Control can not perform zone configuration on this fabric.

Action

Remove the non-standard zones from the fabric before trying to manage its configuration with IBM Spectrum Control.

BTASD1952E Zone set *zoneset name* is active or in activation or deactivation pending status. Deletion of the zone set is not supported for this fabric. After committing the deactivation, this zone set can be deleted in the next zone control session.

Explanation

This fabric does not support the deletion of the active zone set.

Action

Deactivate the zone set before you delete it.

BTASD1953E Renaming of the active zone set is not supported for this fabric. Deactivate zone set *zoneset name* before attempting to rename it.

Explanation

This fabric does not support the renaming of the active zone set.

Action

Deactivate the zone set before attempting to rename it.

BTASD1954E Zoning cannot be done because there is no connection to the SMI-S provider. Reason: *reason*

Explanation

Ensure that there is a proper connection to the SMI-S provider.

Action

Determine the root cause for the failing SMI-S provider connection. Verify that you entered the correct credentials, restart the SMI-S provider or pass through a firewall.

Check the SMI-S provider connection status.

Restart the SMI-S provider.

BTASD2001W Zoning has changed on this fabric since the configuration panel was opened. You may need to run the Fabric discovery/probe again to get the zone information within IBM Spectrum Control synchronized with the fabric. Do you still want to make zoning changes?

Explanation

The current zone information stored by IBM Spectrum Control is outdated.

Action

In order to synchronize IBM Spectrum Control data with the zoning information from the fabric, you may need to run a Fabric Discovery/Probe for the specific switch associated with the Fabric experiencing this issue.

BTASD2002W This alias is not assigned to any zones. This could result in an error when the zone configuration is applied later. Do you want to continue?.

Explanation

The capabilities of this fabric does not support orphaned aliases. Therefore, creating an alias without a parent zone may result in errors if the zone configuration is applied with this zone still an orphan.

Action

It is recommended that you either add the alias to a zone in this alias wizard or add the alias to a zone using the add or edit zone functionalities.

BTASD2003W No exclusive fabric-wide lock is available on the switches for fabric *fabric name*. Other users might be modifying the zoning configuration from outside of IBM Spectrum Control during your zoning operations.

Explanation

The capabilities of this fabric does not support a fabric-wide lock. Therefore, other clients can change the zoning configuration before this transaction is committed.

Action

Ensure that no other user is modifying the zoning configuration while you make changes to the zoning operation.

BTASD2004W This Out Of Band Agent is already defined with the same parameters. Would you like to save it anyway?

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BTASD3001I A probe job for fabric *fabric name* has been submitted. The inactive Zone Definition for this fabric will have the old Zone Definition until the probe job is complete. Wait a few minutes before working with Zone Definition for this fabric.

Explanation

Changing the zone set data source causes the fabric to be probed again using the newly active zone set data source. Therefore, users should not try to edit the zone configuration for this fabric until the probe is done so that they are working with the most up-to-date zone definition.

Action

Wait a few minutes before working with this fabric so the probe on this fabric can complete.

BTATG - UNIX Command Line Interface (CLI) help messages

- [**BTATG0000E** You must have root user authority to run this program.](#)
- [**BTATG0001E** Invalid option '&1'.](#)
- [**BTATG0003E** The format of the GUID is invalid.](#)
- [**BTATG0004I** A GUID already exists on this host. A new GUID will not be created.](#)
- [**BTATG0005I** A GUID entry was not found. The program is generating a new one.](#)
- [**BTATG0006I** A GUID entry was not found.](#)
- [**BTATG0007E** A GUID entry was not created.](#)
- [**BTATG0008E** The GUID entry could not be written.](#)
- [**BTATG0009E** The GUID entry can not be read.](#)
- [**BTATG0011E** When using '&1' you must enter '&3' or use '&3'.](#)
- [**BTATG0012E** The tivguid program encountered an internal error.](#)
- [**BTATG0013E** '&1' return status is '&2'.](#)

BTATG0000E You must have root user authority to run this program.

Explanation

On a UNIX system, only a root user can run this program.

Action

Log in as root user to run the program.

BTATG0001E Invalid option '&1'.

Explanation

The specified option is not valid.

Action

Check the command syntax and retry the command. Use the -HELP option to see usage and syntax information.

BTATG0003E The format of the GUID is invalid.

Explanation

The GUID value specified is not valid.

Action

Check the GUID value and retry the command. Use the -HELP option to see usage and syntax information.

BTATG0004I A GUID already exists on this host. A new GUID will not be created.

Explanation

There is already a GUID on this host. Use the -WRITE -NEW options to create a new GUID.

Action

Check the command syntax and retry the command. Use the -HELP option to see usage and syntax information.

BTATG0005I A GUID entry was not found. The program is generating a new one.

Explanation

There is no GUID on this system. A new one will be created.

BTATG0006I A GUID entry was not found.

Explanation

There is no GUID on this host.

Action

Create a new GUID or use a known GUID to write it on this host.

BTATG0007E A GUID entry was not created.

Explanation

Creation of the GUID failed.

Action

Contact a Tivoli service representative

BTATG0008E The GUID entry could not be written.

Explanation

Creation of the GUID failed.

Action

Contact a Tivoli service representative.

BTATG0009E The GUID entry can not be read.

Explanation

The attempt to read the GUID failed.

Action

Contact a Tivoli service representative.

BTATG0011E When using '&1' you must enter '&3' or use '&3'.

Explanation

The use of -WRITE was not valid.

Action

Check the command syntax and retry the command. Use the -HELP option to see usage and syntax information.

BTATG0012E The tivguid program encountered an internal error.

Explanation

An internal error occurred.

Action

Contact a Tivoli service representative.

BTATG0013E '&1' return status is '&2'.

Explanation

An error occurred.

Action

Contact a Tivoli service representative.

BTAVM

- [BTVMW4001I Connection test to VMware VI Data Source VMware VI Data Source host name FAILED due to VMWareConnectionStatus](#)
- [BTVMW2013E The addition of the data source Name of the data source failed.](#)
- [BTVMW2014W This VMWare VI Data Source is already defined with the same parameters. Would you like to save it anyway?](#)

BTVMW4001I Connection test to VMware VI Data Source *VMware VI Data Source host name* FAILED due to *VMWareConnectionStatus*

Explanation

The connection test to the named VMware VI Data Source failed.

Action

Check the trace logs to find the error reason. The log files reside in the installation subdirectory device/log. If possible correct the error and rerun the connection test to the VMware VI Data Source. If the problem persists, contact IBM Technical Support.

BTVMW2013E The addition of the data source *Name of the data source* failed.

Explanation

The data source may already exist. To modify an existing data source click on the details button.

Action

Make sure the data source name is not already used. If then the problem persists, contact IBM Technical Support.

BTVMW2014W This VMWare VI Data Source is already defined with the same parameters. Would you like to save it anyway?

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BTAZC - Zone control agent messages

- [BTAZC0001E An error occurred while parsing the Zoning Configuration XML for SAN SAN_name.](#)
- [BTAZC0002E Failed to create or update zone set zone_set_name on the GS3 SAN SAN_name.](#)
- [BTAZC0003E Failed to delete zone set zone_set_name on the GS3 SAN SAN_name.](#)
- [BTAZC0004E Unable to start transaction on the SAN SAN_name.](#)
- [BTAZC0005E Unable to commit a transaction on the Brocade SAN SAN_name.](#)
- [BTAZC0006E Failed to deactivate the zone set zone_set_name on SAN SAN_name.](#)
- [BTAZC0007E Failed to activate the zone set zone_set_name on SAN SAN_name.](#)
- [BTAZC0008E Failed to delete the zone set zone_set_name on SAN SAN_name.](#)
- [BTAZC0009E Failed to delete the zone alias zone_alias_name on SAN SAN_name.](#)
- [BTAZC0010E Failed to create the zone alias zone_alias_name on SAN SAN_name.](#)
- [BTAZC0011E Failed to delete the zone zone_name on SAN SAN_name.](#)
- [BTAZC0012E Failed to create the zone zone_name on SAN SAN_name.](#)
- [BTAZC0013E Failed to create the zone set zone_set_name on SAN SAN_name.](#)
- [BTAZC0014E Failed to deactivate the zone set zone_set_name on SAN SAN_name.](#)
- [BTAZC0015E Failed to add the zone zone_name to the zone set zone_set_name on SAN SAN_name.](#)
- [BTAZC0016E One or more nonstandard zone members present in current zoning configuration for SAN SAN_name. Due to that, the Zone Control Agent will not attempt to modify the zoning configuration for the SAN.](#)
- [BTAZC0017E Start transaction for Zone Control failed.](#)
- [BTAZC0018E Rollback error.](#)
- [BTAZC0019E An error occurred while releasing a resource.](#)
- [BTAZC0020E An error occurred while creating a zone set.](#)
- [BTAZC0021E An error occurred while creating a zone.](#)
- [BTAZC0022E An error occurred while creating an alias.](#)
- [BTAZC0023E An error occurred while creating a member.](#)
- [BTAZC0024E An error occurred while adding a zone to a zone set.](#)
- [BTAZC0025E An error occurred while adding a member to a zone.](#)
- [BTAZC0026E An error occurred while adding an alias to a zone.](#)
- [BTAZC0027E An error occurred while adding a member to an alias.](#)
- [BTAZC0028E An error occurred while removing a zone from zone set.](#)
- [BTAZC0029E An error occurred while removing an alias from a zone.](#)
- [BTAZC0030E An error occurred while removing a member from a zone.](#)
- [BTAZC0031E An error occurred while removing a member from an alias.](#)
- [BTAZC0032E An error occurred while deleting a zone member.](#)
- [BTAZC0033E An error occurred while deleting a zone.](#)
- [BTAZC0034E An error occurred while deleting a zone set.](#)
- [BTAZC0035E An error occurred while deleting an alias.](#)
- [BTAZC0036E An error occurred while activating a zone set.](#)
- [BTAZC0037E An error occurred while deactivating a zone set.](#)
- [BTAZC0038E An error occurred while pinging the Zoning Agent.](#)
- [BTAZC0039E An error occurred during transaction commit action.](#)
- [BTAZC0040E An error occurred while closing a session.](#)
- [BTAZC0041E An error occurred while saving the zone information.](#)
- [BTAZC0042E An error occurred during the Get Capabilities command.](#)
- [BTAZC0043E An error occurred sending the zone control command array.](#)

- [BTAZC0044E An error occurred while sending commands to Switch.](#)
- [BTAZC0045E An error occurred: n unsupported zone database.](#)
- [BTAZC0046E A native error occurred: invalid field length.](#)
- [BTAZC0047E A native error occurred: invalid number of members.](#)
- [BTAZC0048E A native error occurred: invalid arguments.](#)
- [BTAZC0049E A native error occurred: null fabric handle.](#)
- [BTAZC0050E An unknown error occurred during Zone control.](#)
- [BTAZC0051E An XML parse error occurred during Zone Control operations.](#)
- [BTAZC0052E Unable to create logical zone definition.](#)
- [BTAZC0053E An error occurred during Zone Control: library not opened.](#)
- [BTAZC0054E Non standard members in the current zone definition.](#)
- [BTAZC0055E A native error occurred: function not supported.](#)
- [BTAZC0056E An error occurred: not connected to the SAN.](#)
- [BTAZC0057E A native error occurred: invalid buffer index.](#)
- [BTAZC0058E A native error occurred during an HBA API call.](#)
- [BTAZC0059E A native error occurred: no memory available.](#)
- [BTAZC0060E A native error occurred: error loading the HBA API.](#)
- [BTAZC0270E An error occurred during an HBA API call.](#)
- [BTAZC0271E An error occurred during an HBA API call: not supported.](#)
- [BTAZC0272E An error occurred during an HBA API call: invalid handle.](#)
- [BTAZC0273E Bad argument with the HBA API.](#)
- [BTAZC0274E An error occurred during an HBA API call: illegal WWN.](#)
- [BTAZC0275E An error occurred during an HBA API call: illegal index.](#)
- [BTAZC0276E Larger buffer required with the HBA API.](#)
- [BTAZC0277E Information has changed since the last call to HBA RefreshInformation.](#)
- [BTAZC0278E SCSI check condition reported with the HBA API.](#)
- [BTAZC0279E HBA error: adapter may be busy or reserved. Retry may be effective.](#)
- [BTAZC0280E HBA API request timed out. Retry may be effective.](#)
- [BTAZC0281E Referenced HBA has been removed or deactivated.](#)
- [BTAZC0282E Extended Link Service reject with the HBA API.](#)
- [BTAZC0283E An error occurred during an HBA API call: invalid LUN.](#)
- [BTAZC0284E An error occurred during an HBA API call: incompatible.](#)
- [BTAZC0285E Ambiguous WWN with HBA API call.](#)
- [BTAZC0286E Local bus error with the HBA API.](#)
- [BTAZC0287E Local target error with the HBA API.](#)
- [BTAZC0288E Local LUN error with the HBA API.](#)
- [BTAZC0289E HBA API error: local SCSI bound.](#)
- [BTAZC0290E HBA API error on FCID target.](#)
- [BTAZC0291E Target node WWN error with the HBA API.](#)
- [BTAZC0292E Target port WWN error with the HBA API.](#)
- [BTAZC0293E Target LUN error with the HBA API.](#)
- [BTAZC0294E Target LUN ID error with the HBA API.](#)
- [BTAZC0295E An HBA API error occurred: no such binding.](#)
- [BTAZC0296E An HBA API error occurred: not a target.](#)
- [BTAZC0297E Unsupported FC4 with HBA API.](#)
- [BTAZC0298E Incapable error with the HBA API.](#)
- [BTAZC0299E An HBA API error occurred: target busy.](#)
- [BTAZC0301E An error occurred during a zoning command.](#)
- [BTAZC0302E An error occurred during zone control operation.](#)
- [BTAZC0303E CIM native error: Not Supported](#)
- [BTAZC0304E CIM native error: Unspecified Error](#)
- [BTAZC0305E CIM native error: Time Out](#)
- [BTAZC0306E CIM native error: Failed](#)
- [BTAZC0307E CIM native error: Invalid Parameter](#)
- [BTAZC0308E CIM native error: Access Denied](#)
- [BTAZC0309E CIM native error: Not Found](#)
- [BTAZC0310E CIM native error: Already Exist](#)
- [BTAZC0311E CIM native error: Insufficient Resources](#)
- [BTAZC0312E CIM native error: Empty Object Invalid](#)
- [BTAZC0313E CIM native error: No Transaction](#)
- [BTAZC0314E CIM native error: Transaction already on](#)
- [BTAZC0315E CIM native error: Cannot Commit Empty Objects](#)
- [BTAZC0316E CIM native error: Zone Database Full](#)
- [BTAZC0317E CIM native error: Too Many Members](#)
- [BTAZC0318E CIM native error: Fabric is busy](#)
- [BTAZC0319E Failed to create the zone zone_name on SAN SAN_name.](#)
- [BTAZC0320E Failed to create the zone alias zone_alias_name on SAN SAN_name.](#)
- [BTAZC0321E Failed to create the zone set zone_set_name on SAN SAN_name.](#)
- [BTAZC0322E Failed to delete the zone zone_name on SAN SAN_name.](#)
- [BTAZC0323E Failed to delete the zone alias zone_alias_name on SAN SAN_name.](#)
- [BTAZC0324E Failed to delete the zone set zone_set_name on SAN SAN_name.](#)
- [BTAZC0325E Failed to delete the zone member zone_member_name on SAN SAN_name.](#)
- [BTAZC0326E Failed to add the zone zone_name to zone set zone_set_name on SAN SAN_name.](#)
- [BTAZC0327E Failed to add the zone member zone_member_name to zone zone_name on SAN SAN_name.](#)
- [BTAZC0328E Failed to add the zone member zone_member_name to zone alias zone_alias_name on SAN SAN_name.](#)
- [BTAZC0329E Failed to add the zone alias zone_alias_name to zone zone_name on SAN SAN_name.](#)
- [BTAZC0330E Failed to remove the zone zone_name from zone set zone_set_name on SAN SAN_name.](#)
- [BTAZC0331E Failed to remove the zone member zone_member_name from zone zone_name on SAN SAN_name.](#)
- [BTAZC0332E Failed to remove the zone member zone_member_name from zone alias zone_alias_name on SAN SAN_name.](#)
- [BTAZC0333E Failed to remove the zone alias zone_alias_name from zone zone_name on SAN SAN_name.](#)

- [BTAZC0334E Failed to create the zone member zone_member_name on SAN SAN_name.](#)
- [BTAZC0335E Failed to activate Zone Set zone_set_name on SAN SAN_name.](#)
- [BTAZC0336E Failed to deactivate Zone Set zone_set_name on SAN SAN_name.](#)
- [BTAZC0337E Failed to enumerate the CIM entity AdminDomain for SAN SAN_name.](#)
- [BTAZC0338E Failed to start a Transaction for SAN SAN_name.](#)
- [BTAZC0339E Failed to commit a Transaction for SAN SAN_name.](#)
- [BTAZC0340E Failed to close the Session for SAN SAN_name.](#)
- [BTAZC0341E Failed to rollback a Transaction for SAN SAN_name.](#)
- [BTAZC0342E This command is not supported for the Fabric Agent.](#)
- [BTAZC0343E CIM native error: Transaction not available](#)
- [BTAZC0344E Zone must be included in ZoneSet. Zone Name.](#)
- [BTAZC0345E CIM error: Save ZoneDB To Switch Failed.](#)
- [BTAZC0346E CIM error: Save ZoneDBInfo Failed.](#)
- [BTAZC0347E CIM error: Zone Database Error.](#)
- [BTAZC0348E CIM error: Transaction Start Failed.](#)
- [BTAZC0349E CIM error: Transaction End Failed.](#)
- [BTAZC0350E CIM error: Transaction Terminate Failed.](#)
- [BTAZC0351E CIM error: Activate ZoneSet Failed.](#)
- [BTAZC0352E CIM error: Deactivate ZoneSet Failed.](#)
- [BTAZC0353E Unable to access the database to persist zoning changes.](#)
- [BTAZC5000I Started Zone Control layer.](#)
- [BTAZC5001I reserve: SAN=SAN_name, Agent=agent_name, Client=client_name, Token=token_ID.](#)
- [BTAZC5002I startTransaction: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5003I commitTransaction: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5004I rollbackTransaction: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5005I setZoneInfo: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5006I sendCommandArray: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5007I release: SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5008I createZoneSet: zoneSetName=zone_set_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5009I createZone: zoneName=zone_name, zoneType=zone_type, zoneSubType=zone_subtype, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5010I createZoneAlias: zoneAliasName=zone_alias_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5011I createZoneMemberSettingData: zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, targetType=target_type, targetName=target_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5012I addZoneToZoneSet: zoneSetName=zone_set_name, zoneName=zone_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5013I addZoneMemberToZone: zoneName=zone_name, zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5014I addZoneAliasToZone: zoneName=zone_name, zoneAliasName=zone_alias_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5015I addZoneMemberToZoneAlias: zoneAliasName=zone_alias_name, zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5016I removeZoneFromZoneSet: zoneSetName=zone_set_name, zoneName=zone_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5017I removeZoneAliasFromZone: zoneName=zone_name, zoneAliasName=zone_alias_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5018I removeZoneMemberFromZone: zoneName=zone_name, zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5019I removeZoneMemberFromZoneAlias: zoneAliasName=zone_alias_name, zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5020I deleteZoneMember: zoneMemberID=zone_member_id, zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5021I deleteZone: zoneName=zone_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5022I deleteZoneSet: zoneSetName=zone_set_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5023I deleteZoneAlias: zoneAliasName=zone_alias_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5024I activateZoneSet: zoneSetName=zone_set_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5025I deactivateZoneSet: zoneSetName=zone_set_name, SAN=SAN_name, Client=client_name, Token=token_ID, result=return_code.](#)
- [BTAZC5026I ping: SAN=SAN_name, Agent=agent_name.](#)
- [BTAZC5027I readCurrentZoneDefinition: SAN=SAN_name, Client=client_name, Token=token_ID result=return_code.](#)

BTAZC0001E An error occurred while parsing the Zoning Configuration XML for SAN SAN_name.

Explanation

The Zone Control Agent failed to parse an XML stream which defines the current Zoning Configuration for the SAN.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support.

Related reference

- [Getting support](#)

- [🔗 Default locations of log files](#)

BTAZC0002E Failed to create or update zone set *zone_set_name* on the GS3 SAN *SAN_name*.

Explanation

The Zone Control Agent was attempting to create or update a zone set on a fabric using GS3. The Zone Control Agent failed to create or update the zone set, with all of its corresponding zones and the zone's members.

Action

Get the GS3 nativelylog information and contact IBM customer support.

Related reference

- [🔗 Getting support](#)
- [🔗 Default locations of log files](#)

BTAZC0003E Failed to delete zone set *zone_set_name* on the GS3 SAN *SAN_name*.

Explanation

The Zone Control Agent was attempting to delete a zone set on a fabric using GS3. The Zone Control Agent failed to delete the zone set, with all of its corresponding zones and the zone's members.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support.

Related reference

- [🔗 Getting support](#)
- [🔗 Default locations of log files](#)

BTAZC0004E Unable to start transaction on the SAN *SAN_name*.

Explanation

The Zone Control Agent was attempting to establish a session with the fabric. A failure occurred during this process.

Action

Get the Brocade nativelylog, Brocade SMI Agent logs, and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [🔗 Getting support](#)
- [🔗 Default locations of log files](#)

BTAZC0005E Unable to commit a transaction on the Brocade SAN *SAN_name*.

Explanation

The Zone Control Agent has sent one or more Zoning Configuration Command requests to the Brocade fabric. The Zone Control Agent was attempting to commit these requests to the Brocade fabric. A failure occurred during this attempt.

Action

Get the Brocade native log, Brocade SMI Agent logs, and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0006E Failed to deactivate the zone set *zone_set_name* on SAN *SAN_name*.

Explanation

The Zone Control Agent was unable to deactivate the active zone set.

Action

Get the Brocade native log, Brocade SMI Agent logs, and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0007E Failed to activate the zone set *zone_set_name* on SAN *SAN_name*.

Explanation

The Zone Control Agent was unable to activate a specific zone set.

Action

Get the Brocade native log, Brocade SMI Agent logs, and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0008E Failed to delete the zone set *zone_set_name* on SAN *SAN_name*.

Explanation

The Zone Control Agent was unable to delete a specific zone set.

Action

Get the Brocade native log, Brocade SMI Agent logs, and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0009E Failed to delete the zone alias *zone_alias_name* on SAN *SAN_name*.

Explanation

The Zone Control Agent was unable to delete a specific zone alias.

Action

Get the Brocade nativelylog, Brocade SMI Agent logs, and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0010E Failed to create the zone alias *zone_alias_name* on SAN *SAN_name*.

Explanation

The Zone Control Agent was unable to create a specific zone alias.

Action

Get the Brocade nativelylog, Brocade SMI Agent logs, and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0011E Failed to delete the zone *zone_name* on SAN *SAN_name*.

Explanation

The Zone Control Agent was unable to delete a specific zone .

Action

Get the Brocade nativelylog, Brocade SMI Agent logs, and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0012E Failed to create the zone *zone_name* on SAN *SAN_name*.

Explanation

The Zone Control Agent was unable to create a specific zone .

Action

Get the Brocade nativelylog, Brocade SMI Agent logs, and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0013E Failed to create the zone set *zone_set_name* on SAN *SAN_name*.

Explanation

The Zone Control Agent was unable to create a specific zone set.

Action

Get the Brocade native log, Brocade SMI Agent logs, and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0014E Failed to deactivate the zone set *zone_set_name* on SAN *SAN_name*.

Explanation

The Zone Control Agent was unable to deactivate the active zone set.

Action

Get the GS3 native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0015E Failed to add the zone *zone_name* to the zone set *zone_set_name* on SAN *SAN_name*.

Explanation

The Zone Control Agent failed to add a zone to a zoneset.

Action

Get the GS3 native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0016E One or more nonstandard zone members present in current zoning configuration for SAN *SAN_name*. Due to that, the Zone Control Agent will not attempt to modify the zoning configuration for the SAN.

Explanation

The Zone Control Agent will not attempt to modify zone configuration on SAN due to presence of one or more non standard zone members.

Action

Replace the non-standard zone members with equivalent standard zone members using the switch management application. Get the GS3 native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0017E Start transaction for Zone Control failed.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0018E Rollback error.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0019E An error occurred while releasing a resource.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0020E An error occurred while creating a zone set.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0021E An error occurred while creating a zone.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0022E An error occurred while creating an alias.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0023E An error occurred while creating a member.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0024E An error occurred while adding a zone to a zone set.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0025E An error occurred while adding a member to a zone.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0026E An error occurred while adding an alias to a zone.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0027E An error occurred while adding a member to an alias.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0028E An error occurred while removing a zone from zone set.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0029E An error occurred while removing an alias from a zone.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0030E An error occurred while removing a member from a zone.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0031E An error occurred while removing a member from an alias.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0032E An error occurred while deleting a zone member

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0033E An error occurred while deleting a zone.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0034E An error occurred while deleting a zone set.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0035E An error occurred while deleting an alias.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0036E An error occurred while activating a zone set.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0037E An error occurred while deactivating a zone set.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0038E An error occurred while pingging the Zoning Agent.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0039E An error occurred during transaction commit action.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0040E An error occurred while closing a session.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0041E An error occurred while saving the zone information.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0042E An error occurred during the Get Capabilities command

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0043E An error occurred sending the zone control command array.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0044E An error occurred while sending commands to Switch.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0045E An error occurred: n unsupported zone database.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0046E A native error occurred: invalid field length.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0047E A native error occurred: invalid number of members.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0048E A native error occurred: invalid arguments.

Explanation

An error occurred during zone control operation. The new zone entity to be created already exists.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0049E A native error occurred: null fabric handle.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0050E An unknown error occurred during Zone control.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0051E An XML parse error occurred during Zone Control operations.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0052E Unable to create logical zone definition.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0053E An error occurred during Zone Control: library not opened.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0054E Non standard members in the current zone definition.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0055E A native error occurred: function not supported.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0056E An error occurred: not connected to the SAN.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0057E A native error occurred: invalid buffer index.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0058E A native error occurred during an HBA API call.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0059E A native error occurred: no memory available.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0060E A native error occurred: error loading the HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0270E An error occurred during an HBA API call.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0271E An error occurred during an HBA API call: not supported.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0272E An error occurred during an HBA API call: invalid handle.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0273E Bad argument with the HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0274E An error occurred during an HBA API call: illegal WWN.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0275E An error occurred during an HBA API call: illegal index.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0276E Larger buffer required with the HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0277E Information has changed since the last call to HBA_RefreshInformation.

Explanation

An HBA API exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0278E SCSI check condition reported with the HBA API.

Explanation

An HBA API exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0279E HBA error: adapter may be busy or reserved. Retry may be effective.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0280E HBA API request timed out. Retry may be effective.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0281E Referenced HBA has been removed or deactivated.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0282E Extended Link Service reject with the HBA API.

Explanation

An HBA API exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0283E An error occurred during an HBA API call: invalid LUN.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0284E An error occurred during an HBA API call: incompatible.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0285E Ambiguous WWN with HBA API call.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0286E Local bus error with the HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0287E Local target error with the HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0288E Local LUN error with the HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0289E HBA API error: local SCSI bound.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0290E HBA API error on FCID target.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0291E Target node WWN error with the HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0292E Target port WWN error with the HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0293E Target LUN error with the HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0294E Target LUN ID error with the HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0295E An HBA API error occurred: no such binding.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0296E An HBA API error occurred: not a target.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0297E Unsupported FC4 with HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0298E Incapable error with the HBA API.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0299E An HBA API error occurred: target busy.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0301E An error occurred during a zoning command.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0302E An error occurred during zone control operation.

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0303E CIM native error: Not Supported

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0304E CIM native error: Unspecified Error

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0305E CIM native error: Time Out

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0306E CIM native error: Failed

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0307E CIM native error: Invalid Parameter

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0308E CIM native error: Access Denied

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0309E CIM native error: Not Found

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0310E CIM native error: Already Exist

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0311E CIM native error: Insufficient Resources

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0312E CIM native error: Empty Object Invalid

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0313E CIM native error: No Transaction

Explanation

The SMI agent reported that there was no open transaction with the SMI agent or fibre channel switch. If the fabric is a Brocade fabric, this may happen if the Zone Control session was open too long before the changes were applied or if a separate Zone Control session completed while this one was in progress.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support. If this is a Brocade fabric, check Network Advisor logs or check the switch logs to see if another Zone Control session occurred around the same time. IBM customer support may instruct you to increase the Zone Control session timeout setting.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0314E CIM native error: Transaction already on

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0315E CIM native error: Cannot Commit Empty Objects

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0316E CIM native error: Zone Database Full

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0317E CIM native error: Too Many Members

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0318E CIM native error: Fabric is busy

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0319E Failed to create the zone *zone_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to create a specific zone.

Action

Get the CIM native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0320E Failed to create the zone alias *zone_alias_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to create a specific zone alias.

Action

Get the CIM native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0321E Failed to create the zone set *zone_set_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to create a specific zone set.

Action

Get the CIM native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0322E Failed to delete the zone *zone_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to delete a specific zone.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0323E Failed to delete the zone alias *zone_alias_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to delete a specific zone alias.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0324E Failed to delete the zone set *zone_set_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to delete a specific zone set.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0325E Failed to delete the zone member *zone_member_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to delete a specific zone member.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0326E Failed to add the zone *zone_name* to zone set *zone_set_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to add a zone to a specific zone set.

Action

Get the CIM native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0327E Failed to add the zone member *zone_member_name* to zone *zone_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to add a zone member to a specific zone.

Action

Get the CIM native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0328E Failed to add the zone member *zone_member_name* to zone alias *zone_alias_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to add a zone member to a specific zone alias.

Action

Get the CIM native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0329E Failed to add the zone alias *zone_alias_name* to zone *zone_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to add a zone alias to a specific zone.

Action

Get the CIM native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0330E Failed to remove the zone *zone_name* from zone set *zone_set_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to remove a zone from a specific zone set.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0331E Failed to remove the zone member *zone_member_name* from zone *zone_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to remove a zone member from a specific zone.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0332E Failed to remove the zone member *zone_member_name* from zone alias *zone_alias_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to remove a zone member from a specific zone alias.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0333E Failed to remove the zone alias *zone_alias_name* from zone *zone_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to remove a zone alias from a specific zone.

Action

Get the CIM native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0334E Failed to create the zone member *zone_member_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to create a zone member.

Action

Get the CIM native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0335E Failed to activate Zone Set *zone_set_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to activate a specific zone set.

Action

This problem may occur for McDATA fabrics where a Zone Set is empty and an attempt was made to activate this empty Zone Set. Check if the Zone Set is empty. If it is not empty, get the CIM native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0336E Failed to deactivate Zone Set *zone_set_name* on SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable deactivate a specific zone set.

Action

Get the CIM native log information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0337E Failed to enumerate the CIM entity AdminDomain for SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to establish the session with the SAN.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0338E Failed to start a Transaction for SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to start a transaction for the SAN.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0339E Failed to commit a Transaction for SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to commit a transaction to the SAN.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0340E Failed to close the Session for SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to close the session for the SAN.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0341E Failed to rollback a Transaction for SAN *SAN_name*.

Explanation

The SMIS Zone Control Agent was unable to rollback a transaction for the SAN.

Action

Get the CIM nativelylog information and the Zone Control service trace information from the trace log and then contact IBM customer support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0342E This command is not supported for the Fabric Agent.

Explanation

The Zone Control Agent does not support this command.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0343E CIM native error: Transaction not available

Explanation

An exception occurred during zone control operation.

Action

Get the Zone Control service trace information from the trace log and contact IBM customer support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0344E Zone must be included in ZoneSet. *Zone Name*.

Explanation

Zone must be in a zoneset.

Action

Use the ADDZONE Command to associate the Zone with a ZoneSet.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0345E CIM error: Save ZoneDB To Switch Failed.

Explanation

The SMIS Zone Control Agent was unable to save ZoneDB to switch.

Action

Get the Zone Control service trace information from the trace log and contact IBM Software support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0346E CIM error: Save ZoneDBInfo Failed.

Explanation

The SMIS Zone Control Agent was unable to save ZoneDB.

Action

Get the Zone Control service trace information from the trace log and contact IBM Software support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0347E CIM error: Zone Database Error.

Explanation

The SMIS Zone Control Agent encountered a zone database error.

Action

Get the Zone Control service trace information from the trace log and contact IBM Software support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0348E CIM error: Transaction Start Failed.

Explanation

The SMIS Zone Control Agent was unable to start the zone control session.

Action

Get the Zone Control service trace information from the trace log and contact IBM Software support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0349E CIM error: Transaction End Failed.

Explanation

The SMIS Zone Control Agent was unable to end the zone control session.

Action

Get the Zone Control service trace information from the trace log and contact IBM Software support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0350E CIM error: Transaction Terminate Failed.

Explanation

The SMIS Zone Control Agent was unable to terminate the zone control session.

Action

Get the Zone Control service trace information from the trace log and contact IBM Software support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0351E CIM error: Activate ZoneSet Failed.

Explanation

The SMIS Zone Control Agent was unable to activate ZoneSet.

Action

Get the Zone Control service trace information from the trace log and contact IBM Software support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0352E CIM error: Deactivate ZoneSet Failed.

Explanation

The SMIS Zone Control Agent was unable to deactivate ZoneSet.

Action

Get the Zone Control service trace information from the trace log and contact IBM Software support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC0353E Unable to access the database to persist zoning changes .

Explanation

The zone control service needs to save zoning changes to the IBM Spectrum Control database. Access to the database tables was denied. This can occur when a probe of the fabric is in progress.

Action

If a probe of the fabric is in progress, wait for the probe to complete and then try the zone control changes again. If the problem continues, get the Zone Control service trace information from the trace log and contact IBM Software Support.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

BTAZC5000I Started Zone Control layer.

Explanation

The Zone Control layer has been started.

BTAZC5001I reserve: SAN=*SAN_name*, Agent=*agent_name*, Client=*client_name*, Token=*token_ID*.

Explanation

The Zone Control reserve API is used to reserve access of the zone control commands for a particular SAN.

BTAZC5002I startTransaction: SAN=*SAN_name*, Client=*client_name*, Token=*token_ID*, result=*return_code*.

Explanation

The Zone Control startTransaction API is used to begin a zone control transaction on the specified SAN.

BTAZC5003I commitTransaction: SAN=*SAN_name*, Client=*client_name*, Token=*token_ID*. result=*return_code*.

Explanation

The Zone Control commitTransaction API is used to apply zone control changes for this transaction on the specified SAN.

BTAZC5004I rollbackTransaction: SAN=*SAN_name*, Client=*client_name*, Token=*token_ID*. result=*return_code*.

Explanation

The Zone Control rollbackTransaction API is used to remove zone control changes for this transaction that were started on the specified SAN.

BTAZC5005I setZoneInfo: SAN=*SAN_name*, Client=*client_name*,
Token=*token_ID*. result=*return_code*.

Explanation

The Zone Control setZoneInfo API is used to apply zone control changes for this transaction on the specified SAN.

BTAZC5006I sendCommandArray: SAN=*SAN_name*, Client=*client_name*,
Token=*token_ID*. result=*return_code*.

Explanation

The Zone Control sendCommandArray API is used to make zone configuration changes on the specified SAN.

BTAZC5007I release: SAN=*SAN_name*, Client=*client_name*,
Token=*token_ID*. result=*return_code*.

Explanation

The Zone Control release API is used to release access of the zone control commands for a particular SAN.

BTAZC5008I createZoneSet: zoneSetName=*zone_set_name*, SAN=*SAN_name*,
Client=*client_name*, Token=*token_ID*. result=*return_code*.

Explanation

The Zone Control createZoneSet API is used to create a zone set with the specified name.

BTAZC5009I createZone: zoneName=*zone_name*, zoneType=*zone_type*,
zoneSubType=*zone_subtype*, SAN=*SAN_name*, Client=*client_name*,
Token=*token_ID*. result=*return_code*.

Explanation

The Zone Control createZone API is used to create a zone with the specified name, type, and subtype.

BTAZC5010I createZoneAlias: zoneAliasName=*zone_alias_name*,
SAN=*SAN_name*, Client=*client_name*, Token=*token_ID*.
result=*return_code*.

Explanation

The Zone Control createZoneAlias API is used to create a zone alias with the specified name.

BTAZC5011I createZoneMemberSettingData:
zoneMemberID=*zone_member_id*, zoneMemberType=*zone_member_type*,
targetType=*target_type*, targetName=*target_name*, SAN=*SAN_name*,
Client=*client_name*, Token=*token_ID*. result=*return_code*.

Explanation

The Zone Control createZoneMemberSettingData API is used to create a zone member with the specified name and type, and add it to the specified zone or zone alias.

BTAZC5012I addZoneToZoneSet: zoneSetName=zone_set_name,
zoneName=zone_name, SAN=SAN_name, Client=client_name,
Token=token_ID. result=return_code.

Explanation

The Zone Control addZoneToZoneSet API is used to add the zone to the specified zone set.

BTAZC5013I addZoneMemberToZone: zoneName=zone_name,
zoneMemberID=zone_member_id, zoneMemberType=zone_member_type,
SAN=SAN_name, Client=client_name, Token=token_ID.
result=return_code.

Explanation

The Zone Control addZoneMemberToZone API is used to add the zone member to the specified zone.

BTAZC5014I addZoneAliasToZone: zoneName=zone_name,
zoneAliasName=zone_alias_name, SAN=SAN_name, Client=client_name,
Token=token_ID. result=return_code.

Explanation

The Zone Control addZoneAliasToZone API is used to add the zone alias to the specified zone.

BTAZC5015I addZoneMemberToZoneAlias: zoneAliasName=zone_alias_name,
zoneMemberID=zone_member_id, zoneMemberType=zone_member_type,
SAN=SAN_name, Client=client_name, Token=token_ID.
result=return_code.

Explanation

The Zone Control addZoneMemberToZoneAlias API is used to add the zone member to the specified zone alias.

BTAZC5016I removeZoneFromZoneSet: zoneSetName=zone_set_name,
zoneName=zone_name, SAN=SAN_name, Client=client_name,
Token=token_ID. result=return_code.

Explanation

The Zone Control removeZoneFromZoneSet API is used to remove the zone from the specified zone set.

BTAZC5017I removeZoneAliasFromZone: zoneName=zone_name,
zoneAliasName=zone_alias_name, SAN=SAN_name, Client=client_name,
Token=token_ID. result=return_code.

Explanation

The Zone Control removeZoneAliasFromZone API is used to remove the zone alias from the specified zone.

BTZC5018I removeZoneMemberFromZone: zoneName=zone_name,
zoneMemberID=zone_member_id, zoneMemberType=zone_member_type,
SAN=SAN_name, Client=client_name, Token=token_ID.
result=return_code.

Explanation

The Zone Control removeZoneMemberFromZone API is used to remove the zone member from the specified zone.

BTZC5019I removeZoneMemberFromZoneAlias:
zoneAliasName=zone_alias_name, zoneMemberID=zone_member_id,
zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name,
Token=token_ID. result=return_code.

Explanation

The Zone Control removeZoneMemberFromZoneAlias API is used to remove the zone member from the specified zone alias.

BTZC5020I deleteZoneMember: zoneMemberID=zone_member_id,
zoneMemberType=zone_member_type, SAN=SAN_name, Client=client_name,
Token=token_ID. result=return_code.

Explanation

The Zone Control deleteZoneMember API is used to delete the specified zone member.

BTZC5021I deleteZone: zoneName=zone_name, SAN=SAN_name,
Client=client_name, Token=token_ID. result=return_code.

Explanation

The Zone Control deleteZone API is used to delete the specified zone.

BTZC5022I deleteZoneSet: zoneSetName=zone_set_name, SAN=SAN_name,
Client=client_name, Token=token_ID. result=return_code.

Explanation

The Zone Control deleteZoneSet API is used to delete the specified zone set.

BTZC5023I deleteZoneAlias: zoneAliasName=zone_alias_name,
SAN=SAN_name, Client=client_name, Token=token_ID.
result=return_code.

Explanation

The Zone Control deleteZoneAlias API is used to delete the specified zone alias.

BTZC5024I activateZoneSet: zoneSetName=zone_set_name,
SAN=SAN_name, Client=client_name, Token=token_ID.

`result=return_code.`

Explanation

The Zone Control activateZoneSet API is used to make the specified zone set the currently active zone set.

BTAZC5025I deactivateZoneSet: zoneSetName=zone_set_name,
SAN=SAN_name, Client=client_name, Token=token_ID.
result=return_code.

Explanation

The Zone Control deactivateZoneSet API is used to deactivate the specified active zone set.

BTAZC5026I ping: SAN=SAN_name, Agent=agent_name,

Explanation

The Zone Control ping API is used to find an agent to send zone control commands to for the specified SAN, and to gather zone capabilities information about the SAN.

BTAZC5027I readCurrentZoneDefinition: SAN=SAN_name,
Client=client_name, Token=token_ID result=return_code.

Explanation

The Zone Control readCurrentZoneDefinition API is used to gather the current zone information directly from the SAN.

BTM - Common Information Model (CIM) agent messages

- [BTM0001E Unable to connect to the SMI-S provider.](#)
- [BTM0002E CIM intrinsic method failure: value.](#)
- [BTM0003E Unable to disconnect from the SMI-S provider.](#)
- [BTM0004E Error getting Host Initiators connected to Target FCPort: value.](#)
- [BTM0005E Error getting Storage System's FCPorts: value.](#)
- [BTM0006E Error encountered while attempting SMI-S provider discovery.](#)
- [BTM0007E value is not a supported protocol for WBEM.](#)
- [BTM0008E Error getting storage systems from SMI-S provider at value, port value.](#)
- [BTM0009E Unable to get CIM_Product instance for this object: name.](#)
- [BTM0010E Unsupported Profile.](#)
- [BTM0011E Error getting Volumes for Storage System: storage system.](#)
- [BTM0012E Error getting paths from Hosts to Volumes for Storage System: name.](#)
- [BTM0013E Error getting detailed information for Storage System: value.](#)
- [BTM0014E Unable to create CIMObjectPath from String: name.](#)
- [BTM0015E Error getting Storage Pools for Storage System: value.](#)
- [BTM0016E Logical subsystems is an IBM-only concept.](#)
- [BTM0017E Error getting Storage System's logical subsystems: value.](#)
- [BTM0018E Error getting Storage System's Disk Groups: value.](#)
- [BTM0019E Error getting Storage Pools for this logical subsystem: name.](#)
- [BTM0020E Error getting Disks for this Storage Pool: value.](#)
- [BTM0021E Error getting Disks for this Disk Group: value.](#)
- [BTM0022E Error getting Volumes for this Storage Pool: value.](#)
- [BTM0023E Error enumerating namespaces.](#)
- [BTM0024E Error getting a specific Storage Pool: value.](#)
- [BTM0025E Unable to connect to SMI-S provider, bad/missing truststore or incorrect truststore password for SMI-S provider at value.](#)
- [BTM0026E Unable to connect to SMI-S provider, cannot find correct certificate in truststore for SMI-S provider at value.](#)
- [BTM0027E Unable to connect to SMI-S provider. Username, password, and/or protocol may be invalid for SMI-S provider at value.](#)
- [BTM0028E Unable to contact SMI-S provider at value. SMI-S provider may not be running.](#)
- [BTM0029E CIMService's hostname or IP is null.](#)
- [BTM0030E CIMService's port is invalid or null.](#)
- [BTM0031E CIMService's protocol is null.](#)
- [BTM0032E CIMAccessParameterSet's Certificate filename is invalid.](#)
- [BTM0033E CIMAccessParameterSet's password is null.](#)
- [BTM0034E CIMAccessParameterSet's username is null.](#)

- [BTM0035E string must be a CIM ObjectPath String for a Storage System.](#)
- [BTM0036E Invalid set of Volumes. Unable to get PathToLUNs.](#)
- [BTM0037E string must be a CIM ObjectPath String for a Storage Pool.](#)
- [BTM0038W Unable to determine Vendor of Storage System: name.](#)
- [BTM0039E Unable to determine the RAID Level of Volume: name.](#)
- [BTM0040W Unable to get CIM Product info for Storage System: name.](#)
- [BTM0041E Unable to get CIM Product info for Storage System. More than one Chassis associated to Storage System: name.](#)
- [BTM0042W Unable to get CIM Product info for Storage System. No CIM Product associated to Storage System's Chassis: value.](#)
- [BTM0043W More than one CIM Product indirectly associated to Storage System: value.](#)
- [BTM0044E Unable to create Volume object: name.](#)
- [BTM0045E Unable to get Host Initiators that can access this volume: value.](#)
- [BTM0046E Host's permission value not recognized: value.](#)
- [BTM0047E Unable to get Disks for this Storage System: name.](#)
- [BTM0048E More than one Disk Group exists for this Disk: value.](#)
- [BTM0049E No Disk Group associated to this Disk: value.](#)
- [BTM0050E Unable to get Disk Group for Disk: value.](#)
- [BTM0051E This Volume is a component of more than one Storage System: value.](#)
- [BTM0052E This Volume is not part of a Storage System: name.](#)
- [BTM0053E Unable to get Volume: value.](#)
- [BTM0054E Please verify that you are running a supported version of a Common Information Model Agent for the storage subsystem.](#)
- [BTM0055E Unable to return systems associated with cluster: cluster.](#)
- [BTM0056E Unable to return Vendor of Cluster: cluster.](#)
- [BTM0057E Error getting Cluster backend controllers: controllers.](#)
- [BTM0058E cluster must be a CIM Object Path String for a Cluster.](#)
- [BTM0059E Backend Controllers not supported for vendor vendor on Cluster cluster.](#)
- [BTM0060E Unable to get Backend Controllers for this Cluster: cluster.](#)
- [BTM0061E Error getting Cluster managed disks: cluster.](#)
- [BTM0062E Unable to get Managed Disks for this Cluster: cluster.](#)
- [BTM0063E Unable to get Backend Controllers for Backend Volume volume on Cluster cluster.](#)
- [BTM0064E Unable to create Managed Disk object: disk](#)
- [BTM0065E Error getting Managed Disk Group Managed Disks: group](#)
- [BTM0066E Backend Volumes not supported for vendor vendor on Cluster cluster.](#)
- [BTM0067E Unable to get Managed Disks for this Managed Disk Group: group.](#)
- [BTM0068E No clusters associated with this Managed Disk Group: group.](#)
- [BTM0069E Error getting Virtual Disk managed disks: disk.](#)
- [BTM0070E volume must be a CIM Object Path String for a volume.](#)
- [BTM0071E Unable to get Managed Disks for this Virtual Disk: volume.](#)
- [BTM0072E No Clusters associated with this Virtual Disk: volume.](#)
- [BTM0073E Error getting Backend Controller managed disks: disk.](#)
- [BTM0074E controller must be a CIM Object Path String for a SCSI Controller.](#)
- [BTM0075E Unable to get Managed Disks for this Backend Controller: controller.](#)
- [BTM0076E No Clusters associated with this Backend Controller: controller.](#)
- [BTM0077E Error getting Cluster managed Disk Groups: cluster.](#)
- [BTM0078E Unable to get Managed Disk Groups for this Cluster: cluster.](#)
- [BTM0079E Unable to create Managed Disk Group Object: group.](#)
- [BTM0080E Error getting Managed Disk Group Virtual Disks: group.](#)
- [BTM0081E Unable to create Virtual Disk object: volume.](#)
- [BTM0082E Unable to get Cluster virtual disks: cluster.](#)
- [BTM0083E Unable to get Virtual Disks for this Cluster: cluster.](#)
- [BTM0084E Error getting Cluster: cluster.](#)
- [BTM0085E Error getting Storage System Type for Computer System: system.](#)
- [BTM0086E Error checking Storage system Level for Computer System: system.](#)
- [BTM0087E Unable to get the Storage System for this volume: volume.](#)
- [BTM0088E volume must be a CIM Object Path String for a Volume.](#)
- [BTM0089E Detected an unsupported level of the Common Information Model agent.](#)
- [BTM0090E Unable to create CIM Object Path String from Class Definition: class.](#)
- [BTM0091E Unable to determine the Privilege for Host Initiator value to access the Volume value.](#)
- [BTM0092W Cannot get Disk Drives for Storage Pool. No Storage Extents found for this Storage Pool: value.](#)
- [BTM0093E No Storage Extents found for this Disk Drive: value.](#)
- [BTM0094E This SMI-S provider version is not supported.](#)
- [BTM0095E This SMI-S provider vendor is not supported.](#)
- [BTM0096E Unable to retrieve LSI SMI-S CIM provider version.](#)
- [BTM0098E Unable to retrieve CIM Object Path for Storage System: storage system from the SMI-S provider.](#)
- [BTM0100E Cannot find unassigned LUNs because the storage pool list is null.](#)
- [BTM0101E Unable to retrieve the Storage System path.](#)
- [BTM0102E Unable to retrieve the Hardware Account path.](#)
- [BTM0103E The Storage System path is null.](#)
- [BTM0104E The Hardware Account path is null.](#)
- [BTM0105E Unable to retrieve the FC Port path.](#)
- [BTM0106E The FC Port path is null.](#)
- [BTM0107E Unable to retrieve the Authorization Service path for Subsystem: value.](#)
- [BTM0108E The Authorization Service path is null for Subsystem: value.](#)
- [BTM0109E There are multiple Authorization Service paths for the Subsystem: value.](#)
- [BTM0110E Unable to retrieve the Main Controller path for Subsystem: value.](#)
- [BTM0111E The Main Controller path is null for Subsystem: value.](#)
- [BTM0112E There are multiple Main Controller paths for the Subsystem: value.](#)
- [BTM0113E Unable to retrieve Clone Controller path for the Subsystem: value Hardware Account: value FC Port: value.](#)
- [BTM0114E The Clone Controller path is null for Subsystem: value Hardware Account: value FC Port: value.](#)
- [BTM0115E Unable to retrieve the Hardware Account for the Clone Controller: value.](#)
- [BTM0116E The Hardware Account for the Clone Controller: value is null.](#)

- [BTM0117E No Hardware Account for the Clone Controller: value.](#)
- [BTM0118E No Access Control Information for the Clone Controller: name.](#)
- [BTM0119E Unable to retrieve the FC Port for the Clone Controller: name.](#)
- [BTM0120E The FC Port for the Clone Controller: name is null.](#)
- [BTM0121E Unable to create a Clone Controller with FC Port: port number Authorization Service: service.](#)
- [BTM0122E Unable to remove Clone Controller: controller name.](#)
- [BTM0123E Unable to Assign Access with Hardware Account: account number Clone Controller: controller Authorization Service: service.](#)
- [BTM0124E Unable to Remove Access with Hardware Account: account number Clone Controller: controller Authorization Service: service.](#)
- [BTM0125E Unable to Attach Volume with Volume volume name Clone Controller: controller.](#)
- [BTM0126E Unable to Detach Volume with Volume volume name Clone Controller: controller.](#)
- [BTM0127E Unable to get Volume, Subsystem, or AuthorizationService path.](#)
- [BTM0128E Unable to assign Volume value to Path \[name, name\] on Subsystem name using Controller name with Authorization Service name.](#)
- [BTM0129E Unable to unassign Volume name to Path \[name, name\] on Subsystem name using Controller name with Authorization Service name.](#)
- [BTM0130E Rolling back value assignments.](#)
- [BTM0131E Rolling back value unassignments.](#)
- [BTM0132E Error getting unassigned LUNs.](#)
- [BTM0133E Error assigning paths.](#)
- [BTM0134E Error unassigning paths.](#)
- [BTM0141E Unable to Attach Volume with Volume Storage Volume for Controller Controller using Controller Configuration Service: Controller Configuration Service.](#)
- [BTM0142E Unable to Detach Volume with Volume Storage Volume for Controller Controller using Controller Configuration Service: Controller Configuration Service.](#)
- [BTM0149E Error calling extrinsic method {0} rc = {1}: Invalid Storage Pool There are multiple Privilege Management Service paths for the Subsystem: Storage Subsystem.](#)
- [BTM0151E There are multiple Privilege Management Service paths for the Subsystem: Storage Subsystem.](#)
- [BTM0152E The Privilege Management Service path is null for Subsystem: Storage Subsystem.](#)
- [BTM0153E There are multiple Controller Configuration Service paths for the Subsystem: Storage Subsystem.](#)
- [BTM0154E The Controller Configuration Service path is null for Subsystem: Storage Subsystem.](#)
- [BTM0155E Unable to assign Volume Storage Volume to Path \[Hardware Account, FC Port\] on Subsystem Storage Subsystem using Controller Controller with Privilege Management Service Privilege Service and Controller Configuration Service Controller Service.](#)
- [BTM0156E Unable to unassign Volume Storage Volume to Path \[Hardware Account, FC Port\] on Subsystem Storage Subsystem using Controller Controller with Privilege Management Service Privilege Service and Controller Configuration Service Controller Service.](#)
- [BTM0157E Unable to retrieve the model volume path.](#)
- [BTM0158E Unable to assign volume for an invalid client request.](#)
- [BTM0159E Unable to unassign volume for an invalid client request.](#)
- [BTM0200E Unable to create Storage Volume of size value in Storage Pool value.](#)
- [BTM0201E Storage Volume of size value not created in Storage Pool value.](#)
- [BTM0202E Unable to retrieve Storage Service for Storage Pool value.](#)
- [BTM0203E Unable to retrieve Storage System for Storage Pool value.](#)
- [BTM0204E Storage Pool used to create the Storage Volume of size value is null.](#)
- [BTM0205E Size used to create the Storage Volume on Storage Pool value is null.](#)
- [BTM0206E Both the Storage Pool and the size to create the Storage Volume are null.](#)
- [BTM0207E Storage Volume identification is null and Storage Volume cannot be located.](#)
- [BTM0208E Storage Volume identification value failed to retrieve Storage Volume.](#)
- [BTM0209E Storage Volume identification value cannot be used to locate a Storage Volume.](#)
- [BTM0210E Storage Volume object is null for Storage System value.](#)
- [BTM0211E Storage System is null for Storage Volume value.](#)
- [BTM0212E Both the Storage System and the Storage Volume object are null.](#)
- [BTM0213E Unable to return the Paths to Storage Volume value on Storage System value.](#)
- [BTM0214E There are no Paths to Storage Volume value on Storage System value.](#)
- [BTM0215E Client connection is null when retrieving Storage Volume identification value.](#)
- [BTM0216E Storage Volume identification is null.](#)
- [BTM0217E Both the Client connection and the Storage Volume identification are null.](#)
- [BTM0218E Unable to retrieve Storage Volume object using Storage Volume identification value.](#)
- [BTM0219E Storage System Type of value is not valid for Storage Volume identification value.](#)
- [BTM0220E Unable to locate Storage Volume object using Storage Volume identification value.](#)
- [BTM0221E Instance of Storage Volume is null.](#)
- [BTM0222E Unable to retrieve Storage Volume identification from Storage Volume instance.](#)
- [BTM0223E Retrieved invalid Storage System name of value from Storage Volume instance.](#)
- [BTM0224E List of Storage Volume objects is invalid.](#)
- [BTM0225E Unable to complete list of Storage Pool objects for Storage Volume value.](#)
- [BTM0226E Unable to complete list of Storage Pool objects without a Storage Volume object.](#)
- [BTM0227E No Storage Pool objects returned for Storage Volume value.](#)
- [BTM0228E Unable to enumerate Storage Pool objects for Storage Volume value.](#)
- [BTM0229E Unable to return Storage Pool objects for Storage Volume value.](#)
- [BTM0230E Unable to generate a list of Storage Pool objects for Storage Volume value.](#)
- [BTM0231E Unable to generate a list of Storage Pool objects without a Storage Volume object.](#)
- [BTM0232E Unable to create Storage Volumes.](#)
- [BTM0233E Unable to select Storage Pools.](#)
- [BTM0234E More than one Storage Service found for Storage System value.](#)
- [BTM0235E Failed to retrieve newly created Storage Volume of size value in Storage Pool value.](#)
- [BTM0236E Storage Volume to be removed is null.](#)
- [BTM0237E Storage Volume value is not removed.](#)
- [BTM0238E Failed to remove Storage Volume value.](#)
- [BTM0239E Unable to retrieve Storage Service for Storage Volume value.](#)
- [BTM0400E Error calling extrinsic method value rc = value: Unsupported method rc.](#)
- [BTM0401E Error calling extrinsic method value rc = value: Unknown error.](#)
- [BTM0402E Error calling extrinsic method value rc = value: Not Supported.](#)
- [BTM0403E Error calling extrinsic method value rc = value: Failed.](#)
- [BTM0404E Error calling extrinsic method value rc = value: Invalid parameter ports.](#)

- [BTM0405E Error calling extrinsic method value rc = value: Invalid controller.](#)
- [BTM0406E Error calling extrinsic method value rc = value: Missing required property within Subject or Target.](#)
- [BTM0407E Error calling extrinsic method value rc = value: Invalid parameter.](#)
- [BTM0408E Error calling extrinsic method value rc = value: Input controller must have AuthorizationView set to FALSE.](#)
- [BTM0409E Error calling extrinsic method value rc = value: Invalid LogicalDevice instance.](#)
- [BTM0410E Error calling extrinsic method value rc = value: Hardware implementation requires null DeviceNumber.](#)
- [BTM0411E Error calling extrinsic method value rc = value: Input size is bigger than the free spaces left in the InPool.](#)
- [BTM0412E Error calling extrinsic method value rc = value: Authorization failure.](#)
- [BTM0413E Error calling extrinsic method value rc = value: Cannot remove device because it is not attached.](#)
- [BTM0414E Error calling extrinsic method value rc = value: Invalid parameter Subject.](#)
- [BTM0415E Error calling extrinsic method value rc = value: Invalid StorageSetting.](#)
- [BTM0416E Error calling extrinsic method value rc = value: Invalid parameter Target.](#)
- [BTM0417E Error calling extrinsic method value rc = value: Input size is invalid, either less than or equal to 0, or is null.](#)
- [BTM0418E Error calling extrinsic method value rc = value: Access is not yet assigned.](#)
- [BTM0419E Error calling extrinsic method {0} rc = {1}: Invalid Storage Pool method rc = return code: Invalid Storage Pool.](#)
- [BTM0420E Error calling extrinsic method value rc = value: The specified Subject and Target are not associated.](#)
- [BTM0421E Error calling extrinsic method value rc = value: Should remove access first.](#)
- [BTM0422E Error calling extrinsic method value rc = value: Should assign access first.](#)
- [BTM0423E Error calling extrinsic method value rc = value: Element type should be 2, meaning Storage Volume. \(The Volume input should be null.\).](#)
- [BTM0424E Error calling extrinsic method value rc = value: Supports single target only.](#)
- [BTM0425E Error calling extrinsic method value rc = value: UserIDType should be PortWWN.](#)
- [BTM0426E Error calling extrinsic method value rc = value: Cannot create a temporary controller.](#)
- [BTM0427E Error calling extrinsic method value rc = value: LogicalDevice instance is already attached to a Host.](#)
- [BTM0428E Error calling extrinsic method value rc = value: Should detach the device first.](#)
- [BTM0429E Error calling extrinsic method value rc = value: UserIDType should be the same as Name, which is PortWWN.](#)
- [BTM0430E Error calling extrinsic method value rc = value: IBMTSESS cannot create view as specified.](#)
- [BTM0431E Error calling extrinsic method value rc = value: Controller processing fails \(Failed to delete temporary controller from repository server.](#)
- [BTM0432E Error calling extrinsic method value rc = value: IBMTSESS cannot attach the device as specified \(ESSCLI cannot create VolumeAccess as specified\).](#)
- [BTM0433E Error calling extrinsic method value rc = value: IBMTSESS cannot remove the device as specified \(ESSCLI cannot delete VolumeAccess as specified\).](#)
- [BTM0434E Error calling extrinsic method value rc = value: IBMTSESS does not support modification of volume.](#)
- [BTM0435E Error calling extrinsic method value rc = value: IBMTSESS cannot AssignAccess as specified \(ESSCLI cannot create or set HostConnection\).](#)
- [BTM0436E Error calling extrinsic method value rc = value: IBMTSESS cannot RemoveAccess as specified \(ESSCLI cannot delete or set HostConnection\).](#)
- [BTM0437E Error calling extrinsic method value rc = value: HardwareAccount instance already exists or HardwareAccount processing fails.](#)
- [BTM0438E Error calling extrinsic method value rc = value: HardwareAccount processing fails.](#)
- [BTM0439E Error calling extrinsic method value rc = value: IBMTSESS cannot create volume as specified \(ESSCLI cannot create volume\).](#)
- [BTM0440E Error calling extrinsic method value rc = value: Controller processing failed.](#)
- [BTM0441E Error calling extrinsic method value rc = value: HardwareAccount processing failed.](#)
- [BTM0442E Error calling extrinsic method name rc = value: Creating indication failure.](#)
- [BTM0443E Error calling extrinsic method name rc = value: The requested logical subsystem already contains the maximum number of volumes allowed.](#)
- [BTM0444E Error calling extrinsic method value rc = value: The requested amount of volume addresses exceeds the maximum number of volumes allowed in the given logical subsystems.](#)
- [BTM0459E Unable to get CIM StorageExtent instance for this object: disk drive from the SMI-S provider.](#)
- [BTM0460E Unable to get CIM PhysicalPackage instance for this object: physical package from the SMI-S provider.](#)
- [BTM0461E Unable to get CIM SoftwareIdentity instance for this object: software identity from the SMI-S provider.](#)
- [BTM0462E Error calling extrinsic method method rc = return code: Invalid Protocol.](#)
- [BTM0463E Error calling extrinsic method method rc = return code: Cannot create temporary controller in SMI-S provider repository.](#)
- [BTM0464E Unable to retrieve CIM SystemSpecificCollection paths for CIM Privilege: CIM Object.](#)
- [BTM0465E Unable to retrieve CIM StorageHardwareID paths for CIM SystemSpecificCollection: CIM Object.](#)
- [BTM0466E Unable to retrieve CIM StorageHardwareID paths for CIM Privilege: CIM Object.](#)
- [BTM0467E Unable to retrieve CIM Privilege paths for CIM SCSIProtocolController: CIM Object.](#)
- [BTM0468E Unable to retrieve CIM SCSIProtocolController paths for CIM StorageVolume: CIM Object.](#)
- [BTM0469E Cannot determine if CIM StorageVolume has been surfaced: CIM Object.](#)
- [BTM0470E Cannot retrieve cache size for CIM ComputerSystem: CIM Object.](#)
- [BTM0550W Cannot get Disk Drives for Storage Pool. No Disk Drives found for this Storage Pool: value.](#)
- [BTM0551W Cannot get Disk Drives for this Storage Pool: value.](#)
- [BTM0552W Value of value not available for: property.](#)
- [BTM0553I Probing Disks for DiskGroup: value](#)
- [BTM0554I Probing Disks for StoragePool: value](#)
- [BTM0555I Number of Disks Found Currently: value. {0}. Continuing to Probe Disks.](#)
- [BTM0556W Cannot get Disk Drives for Storage System. No Disk Drives found for this Storage System: value.](#)
- [BTM0557W Cannot get Disk Drives for this Storage System: value.](#)
- [BTM0558I Number of Volumes Found Currently: value. Continuing to Probe Volumes.](#)
- [BTM0559I Probing Volumes for StoragePool: value.](#)
- [BTM0560I Probing Volumes for Storage System: value.](#)
- [BTM0561I Probing Disks for Storage System: value.](#)
- [BTM0562I Probing Storage Pools for Storage System: value.](#)
- [BTM0563I Probing properties of Storage System: value.](#)
- [BTM0564W More than one CIM Product indirectly associated to the following Device: value.](#)
- [BTM0565W Exception caught while getting CIM Product info for Storage System: value.](#)
- [BTM0566W Unable to get CIM Product info for Device: value.](#)
- [BTM0567W Exception caught while getting CIM Product info for Device: value.](#)
- [BTM0568I value Volumes Found.](#)
- [BTM0569I value Disks Found.](#)
- [BTM0571W Exception caught while getting Host Initiators that can access this volume: value.](#)
- [BTM0572W Exception caught while trying to determine RAID Level for StoragePool: value.](#)
- [BTM0573E Exception caught while formatting this Host Bus Adapter port World Wide Name: value.](#)
- [BTM0574W Capacity of Disk Drive is not available: value.](#)
- [BTM0575W Exception caught while getting Host Initiators access to Volumes through this View: value.](#)

- [BTM0576I Probing Views of Host Initiator access to Volumes.](#)
- [BTM0577I value Views Found.](#)
- [BTM0578E Unable to connect to SMI-S provider. None of the default namespaces are valid for this SMI-S provider.](#)
- [BTM0600E Unable to get Array Site for Disk: value.](#)
- [BTM0601E Error calling extrinsic method value rc = value: A timeout occurred trying to call the method.](#)
- [BTM0602E Error calling extrinsic method value rc = value: The instance of the Logical Device is invalid.](#)
- [BTM0603E Error calling extrinsic method value rc = value: There is a conflict in the Device Number.](#)
- [BTM0604E Error calling extrinsic method value rc = value: A Device Number parameter must be provided.](#)
- [BTM0605E Error calling extrinsic method value rc = value: A null Device Number is required by the device.](#)
- [BTM0606E Error calling extrinsic method value rc = value: The device is busy.](#)
- [BTM0607E Error calling extrinsic method value rc = value: The Protocol Controller is invalid.](#)
- [BTM0608E Error calling extrinsic method value rc = value: The volume types are invalid.](#)
- [BTM0609E Error calling extrinsic method value rc = value: One or more parameters are in the wrong System Scope.](#)
- [BTM0610E Error calling extrinsic method value rc = value: The controller needs to be created first.](#)
- [BTM0611E Error calling extrinsic method value rc = value: The ESSCLI call to create the volume access failed.](#)
- [BTM0612E Error calling extrinsic method value rc = value: The ESSCLI call to list the volume access failed.](#)
- [BTM0613E Error calling extrinsic method value rc = value: The instance of the Logical Device is not associated with the Controller.](#)
- [BTM0614E Error calling extrinsic method value rc = value: The subject is not supported.](#)
- [BTM0615E Error calling extrinsic method value rc = value: The Privilege is not supported.](#)
- [BTM0616E Error calling extrinsic method value rc = value: The Target is not supported.](#)
- [BTM0617E Error calling extrinsic method value rc = value: A null parameter is not supported.](#)
- [BTM0618E Error calling extrinsic method value rc = value: Configuration Service is in use.](#)
- [BTM0619E Error calling extrinsic method value rc = value: The size is invalid.](#)
- [BTM0620E Error calling extrinsic method value rc = value: The Element Type is invalid.](#)
- [BTM0621E Error calling extrinsic method value rc = value: The Goal is invalid.](#)
- [BTM0622E Error calling extrinsic method value rc = value: The Storage Pool is invalid.](#)
- [BTM0623E Error calling extrinsic method value rc = value: The redundancy for the Storage Pool is invalid.](#)
- [BTM0624E Error calling extrinsic method value rc = value: The requested Data Type does not match the Data Type for the Storage Pool.](#)
- [BTM0625E Error calling extrinsic method value rc = value: The Data Type is invalid.](#)
- [BTM0626E Error calling extrinsic method value rc = value: The Element is invalid.](#)
- [BTM0627E Error calling extrinsic method value rc = value: No parameters were specified for the modification.](#)
- [BTM0628E Error calling extrinsic method value rc = value: Unable to create volume.](#)
- [BTM0629E Error calling extrinsic method value rc = value: The LSS already contains the maximum number of volumes.](#)
- [BTM0630E Error calling extrinsic method value rc = value: There are not enough volume addresses in the LSS.](#)
- [BTM0631E Error calling extrinsic method value rc = value: The Identification parameter is missing or not unique.](#)
- [BTM0632E Error calling extrinsic method value rc = value: A null Ports parameter is required by the Controller.](#)
- [BTM0633E Error calling extrinsic method value rc = value: The Controller is busy.](#)
- [BTM0634E Error calling extrinsic method value rc = value: The Identity is invalid.](#)
- [BTM0635E Error calling extrinsic method value rc = value: The Element Name is invalid.](#)
- [BTM0636E Error calling extrinsic method value rc = value: The Protocol is invalid.](#)
- [BTM0637E Error calling extrinsic method value rc = value: The Privilege is invalid.](#)
- [BTM0638E Error calling extrinsic method value rc = value: The Ports are invalid.](#)
- [BTM0639E Error calling extrinsic method value rc = value: The host connection could not be deleted.](#)
- [BTM0640E Error calling extrinsic method value rc = value: The host connection could not be created.](#)
- [BTM0641E Error calling extrinsic method value rc = value: The host connection could not be set.](#)
- [BTM0642E Error calling extrinsic method value rc = value: No Ports are available in this configuration.](#)
- [BTM0701I Probing Managed Disks for Managed Disk Group: value](#)
- [BTM0702I Number of Managed Disks currently found: value. Continuing to probe managed disks.](#)
- [BTM0703I value Managed Disks found.](#)
- [BTM0704I Probing Virtual Disks for Cluster: value](#)
- [BTM0705I Number of Virtual Disks currently found: value. Continuing to probe Virtual Disks.](#)
- [BTM0706I value Virtual Disks found.](#)
- [BTM0707I Probing Virtual Disks for Managed Disk Group: value](#)
- [BTM0708I Probing Managed Disks for Cluster: value](#)
- [BTM0709I Probing Managed Disks for Backend Controller: value](#)
- [BTM0710E Unable to retrieve data for Managed Disk: value](#)
- [BTM0711E Unable to retrieve data for Managed Disk Group: value](#)
- [BTM0712E Unable to retrieve data for Virtual Disk: value](#)
- [BTM0713E Unable to retrieve data for Backend Controller: value](#)
- [BTM0714E Unable to retrieve data for FC Port: value](#)
- [BTM0715E Unable to retrieve data for value Managed Disk\(s\) among the value Managed Disks found.](#)
- [BTM0716E Unable to retrieve data for value Virtual Disk\(s\) among the value Virtual Disks found.](#)
- [BTM0717E Unable to retrieve FC Ports for Cluster: value](#)

BTM0001E Unable to connect to the SMI-S provider.

Explanation

Unable to connect to the SMI-S provider.

Action

Verify the SMI-S provider host name or IP address and port.

BTM0002E CIM intrinsic method failure: *value*.

Explanation

The CIM intrinsic method failure is specified.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0003E Unable to disconnect from the SMI-S provider.

Explanation

Unable to disconnect from the SMI-S provider.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0004E Error getting Host Initiators connected to Target FCPort: *value*.

Explanation

Error getting Host Initiators connected to the specified target FCPort.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0005E Error getting Storage System's FCPorts: *value*.

Explanation

Error getting the specified storage system's FCPorts.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0006E Error encountered while attempting SMI-S provider discovery.

Explanation

Error encountered while attempting SMI-S provider discovery.

Action

Manually enter SMI-S providers in the SMI-S provider Login Administration Panel or contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0007E *value* is not a supported protocol for WBEM.

Explanation

The specified protocol is not a supported protocol for WBEM.

Action

Use http or https.

BTM0008E Error getting storage systems from SMI-S provider at *value*, port *value*.

Explanation

Error getting storage systems from SMI-S provider.

Action

Verify that the SMI-S provider is configured with one or more storage systems. If problem still exists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0009E Unable to get CIM_Product instance for this object: *name*.

Explanation

Unable to get CIM_Product instance for the specified object.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0010E Unsupported Profile.

Explanation

Unsupported Profile.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0011E Error getting Volumes for Storage System: *storage system*.

Explanation

Error getting the specified storage system's volumes.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0012E Error getting paths from Hosts to Volumes for Storage System: *name*.

Explanation

Error getting paths from Hosts to Volumes for the specified storage system.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0013E Error getting detailed information for Storage System: *value*.

Explanation

Error getting detailed information for the specified storage system.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0014E Unable to create CIMObjectPath from String: *name*.

Explanation

Unable to create CIMObjectPath from the specified string.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0015E Error getting Storage Pools for Storage System: *value*.

Explanation

Error getting storage system's specified storage pools.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0016E Logical subsystems is an IBM-only concept.

Explanation

Logical subsystems is an IBM-only concept.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0017E Error getting Storage System's logical subsystems: *value*.

Explanation

Error getting the storage system's specified logical subsystems.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0018E Error getting Storage System's Disk Groups: *value*.

Explanation

Error getting the specified storage system's disk groups.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

BTM0019E Error getting Storage Pools for this logical subsystem: *name*.

Explanation

Error getting storage pools for the specified logical subsystem.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

BTM0020E Error getting Disks for this Storage Pool: *value*.

Explanation

Error getting Disks for the specified storage pool.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

BTM0021E Error getting Disks for this Disk Group: *value*.

Explanation

Error getting Disks for the specified disk group.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

BTM0022E Error getting Volumes for this Storage Pool: *value*.

Explanation

Error getting Volumes for the specified storage pool.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

BTM0023E Error enumerating namespaces.

Explanation

Error enumerating namespaces.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0024E Error getting a specific Storage Pool: *value*.

Explanation

Error getting the specified storage pool.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0025E Unable to connect to SMI-S provider, bad/missing truststore or incorrect truststore password for SMI-S provider at *value*.

Explanation

Unable to connect to SMI-S provider, bad/missing truststore or incorrect truststore password for SMI-S provider.

Action

Verify that the SMI-S provider's correct username, password and truststore are configured in the SMI-S provider Login Administration Panel. Verify that the SMI-S provider is configured with this username and password. If problem still exists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0026E Unable to connect to SMI-S provider, cannot find correct certificate in truststore for SMI-S provider at *value*.

Explanation

Unable to connect to SMI-S provider, cannot find correct certificate in truststore for SMI-S provider.

Action

Verify that the Certificate File configured for this SMI-S provider in the SMI-S provider Login Administration Panel is correct. The Certificate File for an SMI-S provider may be found on the SMI-S provider machine, and is named 'truststore' by default. If problem still exists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0027E Unable to connect to SMI-S provider. Username, password, and/or protocol may be invalid for SMI-S provider at value.

Explanation

Unable to connect to SMI-S provider. Username, password, and/or protocol may be invalid for SMI-S provider.

Action

Verify that the SMI-S provider's correct username, password and protocol are configured in the SMI-S provider Login Administration Panel. If problem still exists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0028E Unable to contact SMI-S provider at value. SMI-S provider may not be running.

Explanation

Unable to contact SMI-S provider. SMI-S provider may not be running.

Action

Check that SMI-S provider's machine is powered on and that the SMI-S provider is running. If problem still exists, contact IBM customer technical support.

BTM0029E CIMService's hostname or IP is null.

Explanation

CIMService's hostname or IP is null.

Action

Check that all hostnames and IP addresses are valid in the SMI-S provider Login Administration Panel.

BTM0030E CIMService's port is invalid or null.

Explanation

CIMService's port is invalid or null.

Action

Check that all ports are valid in the SMI-S provider Login Administration Panel. Typical port values are 5989 and 5988.

BTM0031E CIMService's protocol is null.

Explanation

CIMService's protocol is null.

Action

Check that all protocols are valid in the SMI-S provider Login Administration Panel. Supported protocols are http and https.

BTM0032E CIMAccessParameterSet's Certificate filename is invalid.

Explanation

CIMAccessParameterSet's Certificate filename is invalid.

Action

Check the Certificate Files listed in the SMI-S provider Login Administration Panel. SMI-S providers that use the https protocol require a valid Certificate File from the SMI-S provider. SMI-S providers that use the http protocol do not require a Certificate File.

BTM0033E CIMAccessParameterSet's password is null.

Explanation

CIMAccessParameterSet's password is null.

Action

Check that all passwords are valid in the SMI-S provider Login Administration Panel.

BTM0034E CIMAccessParameterSet's username is null.

Explanation

CIMAccessParameterSet's username is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0035E *string* must be a CIM ObjectPath String for a Storage System.

Explanation

The specified string must be a CIM ObjectPath String for a storage system.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0036E Invalid set of Volumes. Unable to get PathToLUNs.

Explanation

Invalid set of Volumes. Unable to get PathToLUNs.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0037E *string* must be a CIM ObjectPath String for a Storage Pool.

Explanation

The specified string must be a CIM ObjectPath String for a storage pool.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0038W Unable to determine Vendor of Storage System: *name*.

Explanation

Unable to determine the vendor of the specified storage system. The Storage System's Vendor might not be reported.

Action

Contact the Storage System vendor and request any SMI-S 1.0.2 CIM Provider updates which might include updates to the PhysicalPackage component of the CIM Provider.

BTM0039E Unable to determine the RAID Level of Volume: *name*.

Explanation

Unable to determine the RAID Level of the specified volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0040W Unable to get CIM_Product info for Storage System: *name*.

Explanation

Unable to get CIM_Product information for the specified storage system.

Action

Contact the Storage System vendor and request any SMI-S 1.0.2 CIM Provider updates which might include updates to the PhysicalPackage component of the CIM Provider.

BTM0041E Unable to get CIM_Product info for Storage System. More than one Chassis associated to Storage System: *name*.

Explanation

Unable to get CIM_Product information for the specified storage system. More than one chassis associated to the storage system.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0042W Unable to get CIM_Product info for Storage System. No CIM_Product associated to Storage System's Chassis: *value*.

Explanation

Unable to get CIM_Product information for the specified storage system. No CIM_Product associated to storage system's Chassis.

Action

Contact the Storage System vendor and request any SMI-S 1.0.2 CIM Provider updates which might include updates to the PhysicalPackage component of the CIM Provider.

BTM0043W More than one CIM_Product indirectly associated to Storage System: *value*.

Explanation

More than one CIM_Product is associated with the storage system's chassis. Inaccurate Model Number, Serial Number, or Firmware may be displayed for this storage system.

Action

Contact the Storage System vendor and request any SMI-S 1.0.2 CIM Provider updates which might include updates to the PhysicalPackage component of the CIM Provider.

BTM0044E Unable to create Volume object: *name*.

Explanation

Unable to create the specified volume object.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0045E Unable to get Host Initiators that can access this volume: *value*.

Explanation

Unable to get Host Initiators that can access the specified volume.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0046E Host's permission value not recognized: *value*.

Explanation

The host's permission value is not recognized.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0047E Unable to get Disks for this Storage System: *name*.

Explanation

Unable to get Disks for the specified storage system.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0048E More than one Disk Group exists for this Disk: *value*.

Explanation

More than one disk group exists for the specified disk.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0049E No Disk Group associated to this Disk: *value*.

Explanation

No disk group associated to the specified disk.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0050E Unable to get Disk Group for Disk: *value*.

Explanation

Unable to get a disk group for the specified disk.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0051E This Volume is a component of more than one Storage System: *value*.

Explanation

The specified volume is a component of more than one storage system.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0052E This Volume is not part of a Storage System: *name*.

Explanation

The specified volume is not part of a storage system.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0053E Unable to get Volume: *value*.

Explanation

Unable to get the specified volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0054E Please verify that you are running a supported version of a Common Information Model Agent for the storage subsystem.

Explanation

Please contact the IBM Support Center.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0055E Unable to return systems associated with cluster: *cluster*.

Explanation

There are no system nodes associated with the managed cluster.

Action

Check the status of the nodes of the cluster using the SAN Volume Controller console. If the nodes are not connected to the cluster, connect the nodes. If the nodes are connected to the cluster, contact IBM customer support.

Related reference

- [Getting support](#)

BTM0056E Unable to return Vendor of Cluster: *cluster*.

Explanation

The vendor of the managed cluster cannot be returned.

Action

Check the status of the nodes of the cluster using the SAN Volume Controller console. If the nodes are not connected to the cluster, connect the nodes. If the nodes are connected to the cluster, contact IBM customer support.

Related reference

- [Getting support](#)

BTM0057E Error getting Cluster backend controllers: *controllers*.

Explanation

Failed to return the backend controllers of the cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0058E *cluster* must be a CIM Object Path String for a Cluster.

Explanation

The string passed to return cluster information was not a valid CIM Object Path String.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0059E Backend Controllers not supported for vendor *vendor* on Cluster *cluster*.

Explanation

Unable to return Backend Controllers for the Cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0060E Unable to get Backend Controllers for this Cluster: *cluster*.

Explanation

Unable to return Backend Controllers for the Cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0061E Error getting Cluster managed disks: *cluster*.

Explanation

Unable to return managed disks for the Cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0062E Unable to get Managed Disks for this Cluster: *cluster*.

Explanation

Unable to return managed disks for the Cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0063E Unable to get Backend Controllers for Backend Volume *volume* on Cluster *cluster*.

Explanation

Unable to return backend controllers for the Cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0064E Unable to create Managed Disk object: *disk*

Explanation

Unable to create managed disk object.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0065E Error getting Managed Disk Group Managed Disks: *group*

Explanation

Unable to return managed disks from a managed disk group.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0066E Backend Volumes not supported for vendor *vendor* on Cluster *cluster*.

Explanation

Unable to return backend volumes from a cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0067E Unable to get Managed Disks for this Managed Disk Group: *group*.

Explanation

Unable to return managed disks from a managed disk group.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0068E No clusters associated with this Managed Disk Group: *group*.

Explanation

No clusters are associated with the managed disk group.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0069E Error getting Virtual Disk managed disks: *disk*.

Explanation

No virtual disks were returned for the specified managed disk.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0070E *volume* must be a CIM Object Path String for a volume.

Explanation

The passed in string is not a CIM object path string for a volume.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0071E Unable to get Managed Disks for this Virtual Disk: *volume*.

Explanation

Unable to return managed disks for the virtual disk.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0072E No Clusters associated with this Virtual Disk: *volume*.

Explanation

Unable to return clusters for the virtual disk.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0073E Error getting Backend Controller managed disks: *disk*.

Explanation

Unable to return backend controller for the managed disk.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0074E *controller* must be a CIM Object Path String for a SCSI Controller.

Explanation

The passed in string is not a valid CIM Object Path String for SCSI Controller.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0075E Unable to get Managed Disks for this Backend Controller: *controller*.

Explanation

Unable to return managed disks for a Backend Controller.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0076E No Clusters associated with this Backend Controller: *controller*.

Explanation

Unable to return Clusters for a Backend Controller.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0077E Error getting Cluster managed Disk Groups: *cluster*.

Explanation

Unable to return managed disk groups for a cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0078E Unable to get Managed Disk Groups for this Cluster: *cluster*.

Explanation

Unable to return managed disk groups for a cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0079E Unable to create Managed Disk Group Object: *group*.

Explanation

Unable to create a managed disk group object.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0080E Error getting Managed Disk Group Virtual Disks: *group*.

Explanation

Unable to return virtual disks for a managed disk group.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0081E Unable to create Virtual Disk object: *volume*.

Explanation

Unable to create a Virtual Disk Object.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0082E Unable to get Cluster virtual disks: *cluster*.

Explanation

Unable to return virtual disks from a cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0083E Unable to get Virtual Disks for this Cluster: *cluster*.

Explanation

Unable to return virtual disks for this cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0084E Error getting Cluster: *cluster*.

Explanation

Unable to return cluster.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0085E Error getting Storage System Type for Computer System: *system*.

Explanation

Unable to return Storage System Type.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0086E Error checking Storage system Level for Computer System: *system*.

Explanation

Unable to check the storage system level.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0087E Unable to get the Storage System for this volume: *volume*.

Explanation

Unable to return storage system for this volume.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0088E *volume* must be a CIM Object Path String for a Volume.

Explanation

The passed CIM object path string is not for a volume.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0089E Detected an unsupported level of the Common Information Model agent.

Explanation

The SMI-S provider that is being used to process CIM objects is not supported.

Action

If the SMI-S provider is an unsupported level of the Common Information Model agent, update the SMI-S provider to a supported level. If not, contact IBM customer support.

Related reference

- [Getting support](#)

BTM0090E Unable to create CIM Object Path String from Class Definition: *class*.

Explanation

The class definition used to create a CIM object path was not a complete CIM object path.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0091E Unable to determine the Privilege for Host Initiator value to access the Volume value.

Explanation

The CIM Object representing a host does not have an associated CIM Object representing a Privilege.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0092W Cannot get Disk Drives for Storage Pool. No Storage Extents found for this Storage Pool: value.

Explanation

The disk drives associated with this storage pool could not be found. The CIM object representing a storage pool does not have an associated CIM Object which represents a storage extent. This storage system's SMI-S 1.0.2 CIM Provider is probably not implementing the extent mapping or disk drive subprofile.

Action

Contact the storage system vendor and request any SMI-S 1.0.2 CIM Provider updates which may include updates to the extent mapping subprofile or the disk drive subprofile of the SMI-S 1.0.2 CIM Provider.

BTM0093E No Storage Extents found for this Disk Drive: value.

Explanation

The CIM Object representing a Disk Drive does not have an associated CIM Object representing a Storage Extent.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0094E This SMI-S provider version is not supported.

Explanation

An SMI-S provider in the SMI-S provider Login Administration Panel is not supported.

Action

Remove the unsupported SMI-S provider from the SMI-S provider Login Administration Panel. Contact IBM customer support for further help.

Related reference

- [Getting support](#)

BTM0095E This SMI-S provider vendor is not supported.

Explanation

An SMI-S provider in the SMI-S provider Login Administration Panel is not supported.

Action

Remove the unsupported SMI-S provider from the SMI-S provider Login Administration Panel. Contact IBM customer support for further help.

Related reference

- [Getting support](#)

BTM0096E Unable to retrieve LSI SMI-S CIM provider version.

Explanation

Unable to retrieve LSI SMI-S CIM provider version.

Action

Contact IBM customer support for further help.

Related reference

- [Getting support](#)

BTM0098E Unable to retrieve CIM Object Path for Storage System: *storage system* from the SMI-S provider.

Explanation

The storage system that previously was discovered by an SMI-S provider can no longer be found. This could be due to the following conditions:

- The storage system was removed from the SMI-S provider where it had been discovered.
- The storage system was removed from the SMI-S provider where it had been discovered, but it was added to another SMI-S provider that is defined.

Action

One of the following actions can be taken:

- If the storage system should not have been removed from the SMI-S provider where it had been discovered, add the storage system to the SMI-S provider.
- If the storage system does not need to be monitored by IBM Spectrum Control, but the storage system information is to be retained, this error will occur when the storage system is probed until the storage system is added back to the SMI-S provider where it had been discovered.
- If the storage system does not need to be monitored and other storage system information is to be retained, uncheck Monitored against the storage system as it is listed in the Storage Subsystem Administration Panel.
- If the storage system should be removed, remove from the SMI-S provider Logins Panel the SMI-S provider where the storage system had been previously discovered. Add back the SMI-S provider, and discover the storage systems again. Check Monitored against all the storage systems that are to be monitored in the Storage Subsystem Administration Panel.
- If the storage system was defined to another SMI-S provider defined to Spectrum Control, remove from the SMI-S provider Logins Panel the SMI-S provider where the storage system had been previously discovered, and the SMI-S provider where the storage system is defined. Add back the SMI-S provider where the storage system was added and the SMI-S provider where the storage system had been previously discovered. Discover the storage systems again. Check Monitored against all the storage systems that are to be monitored in the Storage Subsystem Administration Panel.

Related reference

- [Getting support](#)

BTM0100E Cannot find unassigned LUNs because the storage pool list is null.

Explanation

Cannot find unassigned LUNs because the storage pool list is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0101E Unable to retrieve the Storage System path.

Explanation

Unable to retrieve the storage system path.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0102E Unable to retrieve the Hardware Account path.

Explanation

Unable to retrieve the Hardware Account path.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0103E The Storage System path is null.

Explanation

The storage system path is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0104E The Hardware Account path is null.

Explanation

The Hardware Account path is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0105E Unable to retrieve the FC Port path.

Explanation

Unable to retrieve the FC Port path.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0106E The FC Port path is null.

Explanation

The FC Port path is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0107E Unable to retrieve the Authorization Service path for Subsystem: *value*.

Explanation

Unable to retrieve the Authorization Service path for the specified subsystem.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0108E The Authorization Service path is null for Subsystem: *value*.

Explanation

The Authorization Service path is null for the specified subsystem.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0109E There are multiple Authorization Service paths for the Subsystem: *value*.

Explanation

There are multiple Authorization Service paths for the specified Subsystem.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0110E Unable to retrieve the Main Controller path for Subsystem: *value*.

Explanation

Unable to retrieve the Main Controller path for the specified subsystem.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0111E The Main Controller path is null for Subsystem: *value*.

Explanation

The Main Controller path is null for the specified Subsystem.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0112E There are multiple Main Controller paths for the Subsystem: *value*.

Explanation

There are multiple Main Controller paths for the specified Subsystem.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0113E Unable to retrieve Clone Controller path for the Subsystem: *value* Hardware Account: *value* FC Port: *value*.

Explanation

Unable to retrieve clone controller path for the specified subsystem, hardware account and FC Port.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0114E The Clone Controller path is null for Subsystem: *value* Hardware Account: *value* FC Port: *value*.

Explanation

The clone controller path is null for the specified subsystem, hardware account and FC Port.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0115E Unable to retrieve the Hardware Account for the Clone Controller: *value*.

Explanation

Unable to retrieve the hardware account for the specified clone controller.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0116E The Hardware Account for the Clone Controller: *value* is null.

Explanation

The Hardware Account for the specified clone controller is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0117E No Hardware Account for the Clone Controller: *value*.

Explanation

No Hardware Account for the specified clone controller.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0118E No Access Control Information for the Clone Controller: *name*.

Explanation

No Access Control Information for the specified clone controller.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0119E Unable to retrieve the FC Port for the Clone Controller: *name*.

Explanation

Unable to retrieve the FC Port for the clone controller.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0120E The FC Port for the Clone Controller: *name* is null.

Explanation

The FC Port for the clone controller is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0121E Unable to create a Clone Controller with FC Port: *port number* Authorization Service: *service*.

Explanation

The logical view of LUN Mapping through a specific FC Port was not created in the SMI-S provider.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0122E Unable to remove Clone Controller: *controller name*.

Explanation

The logical view of LUN Mapping through a specific FC Port was not removed from the SMI-S provider.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0123E Unable to Assign Access with Hardware Account: *account number* Clone Controller: *controller* Authorization Service: *service*.

Explanation

Failed to assign access (unmask) for a specific host or HBA to a specific volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0124E Unable to Remove Access with Hardware Account: *account number* Clone Controller: *controller* Authorization Service: *service*.

Explanation

Failed to remove access (mask) for a specific host or HBA to a specific volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0125E Unable to Attach Volume with Volume *volume name* Clone Controller: *controller*.

Explanation

The volume could not be associated with the logical view of the LUN Mapping in the SMI-S provider.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0126E Unable to Detach Volume with Volume *volume name* Clone Controller: *controller*.

Explanation

The volume could not be unassociated with the logical view of the LUN Mapping in the SMI-S provider.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0127E Unable to get Volume, Subsystem, or AuthorizationService path.

Explanation

Unable to get Volume, Subsystem, or AuthorizationService path.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0128E Unable to assign Volume *value* to Path [*name*, *name*] on Subsystem *name* using Controller *name* with Authorization Service

name .

Explanation

Unable to assign the volume to path on the subsystem using the controller with the authorization service.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0129E Unable to unassign Volume *name* to Path [*name*, *name*] on Subsystem *name* using Controller *name* with Authorization Service *name* .

Explanation

Unable to unassign the volume to path on the subsystem using the controller with the authorization service.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0130E Rolling back *value* assignments .

Explanation

Rolling back assignments.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0131E Rolling back *value* unassignments .

Explanation

Rolling back unassignments.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0132E Error getting unassigned LUNs .

Explanation

Error getting unassigned LUNs.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0133E Error assigning paths.

Explanation

Error assigning paths.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0134E Error unassigning paths.

Explanation

Error unassigning paths.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0141E Unable to Attach Volume with Volume *Storage Volume* for Controller *Controller* using Controller Configuration Service: *Controller Configuration Service*.

Explanation

Unable to detach Storage Volume for Controller using Controller Configuration Service.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0142E Unable to Detach Volume with Volume *Storage Volume* for Controller *Controller* using Controller Configuration Service: *Controller Configuration Service*.

Explanation

Unable to detach Storage Volume for Controller using Controller Configuration Service.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0149E Error calling extrinsic method {0} rc = {1}: Invalid Storage Pool There are multiple Privilege Management Service paths for the Subsystem: *Storage Subsystem*.

Explanation

There are multiple Privilege Management Service paths for the Storage Subsystem.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0151E There are multiple Privilege Management Service paths for the Subsystem: *Storage Subsystem*.

Explanation

There are multiple Privilege Management Service paths for the Storage Subsystem.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0152E The Privilege Management Service path is null for Subsystem: *Storage Subsystem*.

Explanation

The Privilege Management Service path is null for the Storage Subsystem.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0153E There are multiple Controller Configuration Service paths for the Subsystem: *Storage Subsystem*.

Explanation

There are multiple Controller Configuration Service paths for the Storage Subsystem.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0154E The Controller Configuration Service path is null for Subsystem: *Storage Subsystem*.

Explanation

The Controller Configuration Service path is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0155E Unable to assign Volume *Storage Volume* to Path [*Hardware Account, FC Port*] on Subsystem *Storage Subsystem* using Controller *Controller* with Privilege Management Service *Privilege Service* and Controller Configuration Service *Controller Service*.

Explanation

Unable to assign Storage Volume to Path [Hardware Account, FC Port] on Storage Subsystem using Controller and Privilege Management Service and Controller Service.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0156E Unable to unassign Volume *Storage Volume* to Path [*Hardware Account, FC Port*] on Subsystem *Storage Subsystem* using Controller *Controller* with Privilege Management Service *Privilege Service* and Controller Configuration Service *Controller Service*.

Explanation

Unable to unassign Storage Volume to Path [Hardware Account, FC Port] on Storage Subsystem using Controller and Privilege Management Service and Controller Service.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0157E Unable to retrieve the model volume path.

Explanation

Unable to retrieve the CIM Object Path for the Storage Volume representing the model volume.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0158E Unable to assign volume for an invalid client request.

Explanation

The request to assign a volume fails because the caller of the request is either invalid, or not supported.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0159E Unable to unassign volume for an invalid client request.

Explanation

The request to unassign a volume fails because the caller of the request is either invalid, or not supported.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0200E Unable to create Storage Volume of size *value* in Storage Pool *value*.

Explanation

Unable to create a storage volume of the specified size in the specified storage pool.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0201E Storage Volume of size *value* not created in Storage Pool *value*.

Explanation

Storage volume of the specified size not created in the specified storage pool.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0202E Unable to retrieve Storage Service for Storage Pool *value*.

Explanation

Unable to retrieve Storage Service for the specified storage pool.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0203E Unable to retrieve Storage System for Storage Pool *value*.

Explanation

Unable to retrieve storage system for the specified storage pool.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0204E Storage Pool used to create the Storage Volume of size *value* is null.

Explanation

Storage pool used to create the storage volume of the specified size is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0205E Size used to create the Storage Volume on Storage Pool *value* is null.

Explanation

Size used to create the storage volume on the specified storage pool is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0206E Both the Storage Pool and the size to create the Storage Volume are null.

Explanation

Both the storage pool and the size to create the storage volume are null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0207E Storage Volume identification is null and Storage Volume cannot be located.

Explanation

Storage volume identification is null and storage volume cannot be located.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0208E Storage Volume identification *value* failed to retrieve Storage Volume.

Explanation

The specified storage volume identification failed to retrieve the specified storage volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0209E Storage Volume identification value cannot be used to locate a Storage Volume.

Explanation

The specified storage volume identification cannot be used to locate a storage volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0210E Storage Volume object is null for Storage System value.

Explanation

Storage volume object is null for the specified storage system.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0211E Storage System is null for Storage Volume value.

Explanation

Storage system is null for the specified storage volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0212E Both the Storage System and the Storage Volume object are null.

Explanation

Both the storage system and the storage volume object are null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0213E Unable to return the Paths to Storage Volume *value* on Storage System *value*.

Explanation

Unable to return the Paths to the specified storage volume on the specified storage system.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0214E There are no Paths to Storage Volume *value* on Storage System *value*.

Explanation

There are no Paths to the specified storage volume on the specified storage system.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0215E Client connection is null when retrieving Storage Volume identification *value*.

Explanation

Client connection is null when retrieving the specified storage volume identification.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0216E Storage Volume identification is null.

Explanation

Storage volume identification is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0217E Both the Client connection and the Storage Volume identification are null.

Explanation

Both the Client connection and the storage volume identification are null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0218E Unable to retrieve Storage Volume object using Storage Volume identification *value*.

Explanation

Unable to retrieve storage volume object using the specified storage volume identification.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0219E Storage System Type of *value* is not valid for Storage Volume identification *value*.

Explanation

The specified storage system type is not valid for the specified storage volume identification.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0220E Unable to locate Storage Volume object using Storage Volume identification *value*.

Explanation

Unable to locate storage volume object using the specified storage volume identification.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0221E Instance of Storage Volume is null.

Explanation

Instance of storage volume is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0222E Unable to retrieve Storage Volume identification from Storage Volume instance.

Explanation

Unable to retrieve storage volume identification from storage volume instance.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0223E Retrieved invalid Storage System name of *value* from Storage Volume instance.

Explanation

Retrieved an invalid storage system name from the storage volume instance.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0224E List of Storage Volume objects is invalid.

Explanation

List of storage volume objects is invalid.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0225E Unable to complete list of Storage Pool objects for Storage Volume *value*.

Explanation

Unable to complete list of storage pool objects for the specified storage volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0226E Unable to complete list of Storage Pool objects without a Storage Volume object.

Explanation

Unable to complete list of storage pool objects without a storage volume object.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0227E No Storage Pool objects returned for Storage Volume *value*.

Explanation

No storage pool objects returned for the specified storage volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0228E Unable to enumerate Storage Pool objects for Storage Volume *value*.

Explanation

Unable to enumerate storage pool objects for the specified storage volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0229E Unable to return Storage Pool objects for Storage Volume *value*.

Explanation

Unable to return storage pool objects for the specified storage volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0230E Unable to generate a list of Storage Pool objects for Storage Volume *value*.

Explanation

Unable to generate a list of storage pool objects for the specified storage volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0231E Unable to generate a list of Storage Pool objects without a Storage Volume object.

Explanation

Unable to generate a list of storage pool objects without a storage volume object.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0232E Unable to create Storage Volumes.

Explanation

Unable to create storage volumes. A previous failure occurred to prevent the creation of a storage volume with a specific size.

Action

Check prior messages for the cause of the failure. If the failure persists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0233E Unable to select Storage Pools.

Explanation

Unable to select storage pools. The storage pools that were derived from the model LUNs did not have sufficient space to satisfy the request. Therefore, storage volume creation would not be possible.

Action

Check if the storage pools that were derived from the model LUNs had enough available space and volume addresses. If there is sufficient space and available volume addresses within each storage pool to create the storage volume, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0234E More than one Storage Service found for Storage System *value*.

Explanation

More than one Storage Service found for the specified storage system.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0235E Failed to retrieve newly created Storage Volume of size *value* in Storage Pool *value*.

Explanation

Failed to retrieve a newly created storage volume of the specified size in the specified storage pool.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0236E Storage Volume to be removed is null.

Explanation

The Storage Volume to be removed does not exist.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0237E Storage Volume *value* is not removed.

Explanation

The Storage Volume is not removed due to an error from the SMI-S provider method.

Action

Contact IBM customer technical support with all related errors.

Related reference

-  [Getting support](#)

BTM0238E Failed to remove Storage Volume *value*.

Explanation

The attempt to remove the Storage Volume failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

-  [Getting support](#)

BTM0239E Unable to retrieve Storage Service for Storage Volume *value*.

Explanation

The attempt to retrieve the Storage Service for the Storage Volume failed.

Action

Contact IBM customer technical support with all related errors.

Related reference

-  [Getting support](#)

BTM0400E Error calling extrinsic method *value* rc = *value*: Unsupported method rc.

Explanation

Error calling the specified extrinsic method. Unsupported method rc.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0401E Error calling extrinsic method *value rc = value*: Unknown error.

Explanation

Error calling the specified extrinsic method. Unknown error.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0402E Error calling extrinsic method *value rc = value*: Not Supported.

Explanation

Error calling the specified extrinsic method. Not Supported.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0403E Error calling extrinsic method *value rc = value*: Failed.

Explanation

Error calling the specified extrinsic method. Failed.

Action

Check the cimom.log file where the SMI-S provider was installed for any messages related to the error occurring within the same time frame. Contact IBM customer technical support with the information gathered from the cimom.log as well as this message.

Related reference

- [Getting support](#)

BTM0404E Error calling extrinsic method *value rc = value*: Invalid parameter ports.

Explanation

Error calling the specified extrinsic method. Invalid parameter ports.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0405E Error calling extrinsic method *value rc = value*: Invalid controller.

Explanation

Error calling the specified extrinsic method. Invalid controller.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0406E Error calling extrinsic method *value rc = value*: Missing required property within Subject or Target.

Explanation

Error calling the specified extrinsic method. Missing required property within Subject or Target.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0407E Error calling extrinsic method *value rc = value*: Invalid parameter.

Explanation

Error calling the specified extrinsic method. Invalid parameter.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0408E Error calling extrinsic method *value rc = value*: Input controller must have AuthorizationView set to FALSE.

Explanation

Error calling the specified extrinsic method. Input controller must have AuthorizationView set to FALSE.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0409E Error calling extrinsic method *value rc = value*: Invalid LogicalDevice instance.

Explanation

Error calling the specified extrinsic method. Invalid LogicalDevice instance.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0410E Error calling extrinsic method *value rc = value*: Hardware implementation requires null DeviceNumber.

Explanation

Error calling the specified extrinsic method. Hardware implementation requires null DeviceNumber.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0411E Error calling extrinsic method *value rc = value*: Input size is bigger than the free spaces left in the InPool.

Explanation

Error calling the specified extrinsic method. Input size is bigger than the free spaces left in the InPool.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0412E Error calling extrinsic method *value rc = value*: Authorization failure.

Explanation

Error calling the specified extrinsic method. Authorization failure.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0413E Error calling extrinsic method *value rc = value*: Cannot remove device because it is not attached.

Explanation

Error calling the specified extrinsic method. Cannot remove device because it is not attached.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0414E Error calling extrinsic method *value rc = value*: Invalid parameter Subject.

Explanation

Error calling the specified extrinsic method. Invalid parameter Subject.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0415E Error calling extrinsic method *value rc = value*: Invalid StorageSetting.

Explanation

Error calling the specified extrinsic method. Invalid StorageSetting.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0416E Error calling extrinsic method *value rc = value*: Invalid parameter Target.

Explanation

Error calling the specified extrinsic method. Invalid parameter Target.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0417E Error calling extrinsic method *value rc = value*: Input size is invalid, either less than or equal to 0, or is null.

Explanation

Error calling the specified extrinsic method. Input size is invalid, either less than or equal to 0, or is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0418E Error calling extrinsic method *value rc = value*: Access is not yet assigned.

Explanation

Error calling the specified extrinsic method. Access is not yet assigned.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0419E Error calling extrinsic method {0} *rc = {1}*: Invalid Storage Pool *method rc = return code*: Invalid Storage Pool.

Explanation

Error calling the specified extrinsic method. Invalid Protocol.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0420E Error calling extrinsic method *value rc = value*: The specified Subject and Target are not associated.

Explanation

Error calling the specified extrinsic method. The specified Subject and Target are not associated.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0421E Error calling extrinsic method *value rc = value*: Should remove access first.

Explanation

Error calling the specified extrinsic method. Should remove access first.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0422E Error calling extrinsic method *value rc = value*: Should assign access first.

Explanation

Error calling the specified extrinsic method. Should assign access first.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0423E Error calling extrinsic method *value rc = value*: Element type should be 2, meaning Storage Volume. (The Volume input should be null.)

Explanation

Error calling the specified extrinsic method. Element type should be 2, meaning storage volume. The volume input should be null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0424E Error calling extrinsic method *value rc = value*: Supports single target only.

Explanation

Error calling the specified extrinsic method. Supports single target only.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0425E Error calling extrinsic method *value rc = value*: UserIDType should be PortWWN.

Explanation

Error calling the specified extrinsic method. UserIDType should be PortWWN.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0426E Error calling extrinsic method *value rc = value*: Cannot create a temporary controller.

Explanation

Error calling the specified extrinsic method. Cannot create a temporary controller.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0427E Error calling extrinsic method *value rc = value*: LogicalDevice instance is already attached to a Host.

Explanation

Error calling the specified extrinsic method. LogicalDevice instance is already attached to a Host.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0428E Error calling extrinsic method *value rc = value*: Should detach the device first.

Explanation

Error calling the specified extrinsic method. Should detach the device first.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0429E Error calling extrinsic method *value rc = value*: UserIDType should be the same as Name, which is PortWWN.

Explanation

Error calling the specified extrinsic method. UserIDType should be the same as Name, which is PortWWN.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0430E Error calling extrinsic method *value rc = value*: IBMTSESS cannot create view as specified.

Explanation

Error calling the specified extrinsic method. IBMTSESS cannot create view as specified.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0431E Error calling extrinsic method *value rc = value*: Controller processing fails (Failed to delete temporary controller from repository server).

Explanation

Error calling the specified extrinsic method. Controller processing fails (Failed to delete temporary controller from repository server).

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0432E Error calling extrinsic method value rc = value: IBMTSESS cannot attach the device as specified (ESSCLI cannot create VolumeAccess as specified).

Explanation

Error calling the specified extrinsic method. IBMTSESS cannot attach the device as specified (ESSCLI cannot create VolumeAccess as specified).

Action

Check the cimom.log file where the SMI-S provider was installed for any ESSCLI messages related to the error occurring within the same time frame. If necessary, contact IBM customer technical support with the information gathered from the cimom.log as well as this message.

Related reference

- [Getting support](#)

BTM0433E Error calling extrinsic method value rc = value: IBMTSESS cannot remove the device as specified (ESSCLI cannot delete VolumeAccess as specified).

Explanation

Error calling the specified extrinsic method. IBMTSESS cannot remove the device as specified (ESSCLI cannot delete VolumeAccess as specified).

Action

Check the cimom.log file where the SMI-S provider was installed for any ESSCLI messages related to the error occurring within the same time frame. If necessary, contact IBM customer technical support with the information gathered from the cimom.log as well as this message.

Related reference

- [Getting support](#)

BTM0434E Error calling extrinsic method value rc = value: IBMTSESS does not support modification of volume.

Explanation

Error calling the specified extrinsic method. IBMTSESS does not support modification of volume.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0435E Error calling extrinsic method value rc = value: IBMTSESS cannot AssignAccess as specified (ESSCLI cannot create or set HostConnection).

Explanation

Error calling the specified extrinsic method. IBMTSESS cannot AssignAccess as specified (ESSCLI cannot create or set HostConnection).

Action

Check the cimom.log file where the SMI-S provider was installed for any ESSCLI messages related to the error occurring within the same time frame. If necessary, contact IBM customer technical support with the information gathered from the cimom.log as well as this message.

Related reference

- [Getting support](#)

BTM0436E Error calling extrinsic method *value rc = value: IBMTSESS cannot RemoveAccess as specified (ESSCLI cannot delete or set HostConnection)*.

Explanation

Error calling the specified extrinsic method. IBMTSESS cannot RemoveAccess as specified (ESSCLI cannot delete or set HostConnection).

Action

Check the cimom.log file where the SMI-S provider was installed for any ESSCLI messages related to the error occurring within the same time frame. If necessary, contact IBM customer technical support with the information gathered from the cimom.log as well as this message.

Related reference

- [Getting support](#)

BTM0437E Error calling extrinsic method *value rc = value: HardwareAccount instance already exists or HardwareAccount processing fails*.

Explanation

Error calling the specified extrinsic method. HardwareAccount instance already exists or HardwareAccount processing fails.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0438E Error calling extrinsic method *value rc = value: HardwareAccount processing fails*.

Explanation

Error calling the specified extrinsic method. HardwareAccount processing fails.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0439E Error calling extrinsic method *value rc = value: IBMTSESS cannot create volume as specified (ESSCLI cannot create volume)*.

Explanation

Error calling the specified extrinsic method. IBMTSESS cannot create volume as specified (ESSCLI cannot create volume).

Action

Check the cimom.log file where the SMI-S provider was installed for any ESSCLI messages related to the error occurring within the same time frame. If necessary, contact IBM customer technical support with the information gathered from the cimom.log as well as this message.

Related reference

-  [Getting support](#)

BTM0440E Error calling extrinsic method *value rc = value: Controller processing failed*.

Explanation

Error calling the specified extrinsic method. Controller processing failed.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0441E Error calling extrinsic method *value rc = value: HardwareAccount processing failed*.

Explanation

Error calling the specified extrinsic method. HardwareAccount processing failed.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0442E Error calling extrinsic method *name rc = value: Creating indication failure*.

Explanation

Error calling the specified extrinsic method. Creating indication failure.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0443E Error calling extrinsic method *name* rc = *value*: The requested logical subsystem already contains the maximum number of volumes allowed.

Explanation

Error calling the specified extrinsic method. The requested logical subsystem already contains the maximum number of volumes allowed. This problem occurs when there is enough volume addresses within the logical subsystem to satisfy the request, but the requested volume size is beyond the capacity allowed by the logical subsystem.

Action

If the model LUNs used to provision a file system were all exhausted, create a policy where the model LUNs outside the volume group are used. If that is still not sufficient, assign a model LUN to the host from a logical subsystem with sufficient space. If that does not resolve the problem, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0444E Error calling extrinsic method *value* rc = *value*: The requested amount of volume addresses exceeds the maximum number of volumes allowed in the given logical subsystems.

Explanation

Error calling the specified extrinsic method. The requested amount of volume addresses exceeds the maximum number of volumes allowed in the given logical subsystems. This problem occurs when there are not enough volume addresses within the logical subsystem to satisfy the request. If the volume being created will be the last available volume address within the logical subsystem, it must be the same size of the remaining free space within the logical subsystem. Otherwise, this error occurs.

Action

If the model LUNs used to provision a file system were all exhausted, create a policy where the model LUNs outside the volume group are used. If that is still not sufficient, assign a model LUN to the host from a logical subsystem with sufficient space and available volume addresses. If that does not resolve the problem, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0459E Unable to get CIM_StorageExtent instance for this object: *disk drive* from the SMI-S provider.

Explanation

Cannot retrieve the storage extent for this disk drive object from the SMI-S provider. The probe of the storage system will abort.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0460E Unable to get CIM_PhysicalPackage instance for this object: *physical package* from the SMI-S provider.

Explanation

Cannot retrieve the physical package for this disk drive object from the SMI-S provider. The probe of the storage system will abort.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0461E Unable to get CIM_SoftwareIdentity instance for this object: *software identity* from the SMI-S provider.

Explanation

Cannot retrieve the software identity for this disk drive object from the SMI-S provider. The probe of the storage system will abort.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0462E Error calling extrinsic method *method rc = return code*: Invalid Protocol.

Explanation

Error calling the specified extrinsic method. Invalid Protocol.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0463E Error calling extrinsic method *method rc = return code*: Cannot create temporary controller in SMI-S provider repository.

Explanation

Error calling the specified extrinsic method.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0464E Unable to retrieve CIM_SystemSpecificCollection paths for CIM_Privilege: *CIM Object*.

Explanation

The CIM Object representing a System Specific Collection does not have an associated CIM Object representing a Privilege.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0465E Unable to retrieve CIM_StorageHardwareID paths for CIM_SystemSpecificCollection: *CIM Object*.

Explanation

The CIM Object representing a Storage Hardware ID does not have an associated CIM Object representing a System Specific Collection.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0466E Unable to retrieve CIM_StorageHardwareID paths for CIM_Privilege: *CIM Object*.

Explanation

The CIM Object representing a Storage Hardware ID does not have an associated CIM Object representing a Privilege.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0467E Unable to retrieve CIM_Privilege paths for CIM_SCSIProtocolController: *CIM Object*.

Explanation

The CIM Object representing a Privilege does not have an associated CIM Object representing a Controller.

Action

Contact IBM customer support.

Related reference

-  [Getting support](#)

BTM0468E Unable to retrieve CIM_SCSIProtocolController paths for CIM_StorageVolume: *CIM Object*.

Explanation

The CIM Object representing a Controller does not have an associated CIM Object representing a Storage Volume.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0469E Cannot determine if CIM_StorageVolume has been surfaced: CIM Object.

Explanation

Cannot determine if the CIM Object representing a Storage Volume has been surfaced on a host.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0470E Cannot retrieve cache size for CIM_ComputerSystem: CIM Object.

Explanation

The cache size of the Storage System cannot be determined from the CIM Object representing a Storage System.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

BTM0550W Cannot get Disk Drives for Storage Pool. No Disk Drives found for this Storage Pool: value.

Explanation

The disk drives associated with the specified storage pool could not be found. This storage system's SMI-S 1.0.2 CIM provider is probably not implementing the extent mapping or disk drive subprofile.

Action

Contact the storage system vendor and request any SMI-S 1.0.2 CIM Provider updates which might include updates to the extent mapping subprofile or the disk drive subprofile of the SMI-S 1.0.2 CIM Provider.

BTM0551W Cannot get Disk Drives for this Storage Pool: value.

Explanation

An error occurred while retrieving disk drives associated to this storage pool.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0552W Value of *value* not available for: *property* .

Explanation

The value of a CIM property that is expected from the storage system's SMI-S 1.0.2 CIM Provider cannot be determined.

Action

Contact the Storage System vendor and request any SMI-S 1.0.2 CIM Provider updates which might include an update that provides this property's value.

BTM0553I Probing Disks for DiskGroup: *value*

Explanation

The probe is finding the disks for this DiskGroup.

BTM0554I Probing Disks for StoragePool: *value*

Explanation

The probe is finding the disks for this storage pool.

BTM0555I Number of Disks Found Currently: *value*. {0}. Continuing to Probe Disks.

Explanation

The probe is finding the disks. This status update is to inform how many disks have been processed at this point during the probe.

BTM0556W Cannot get Disk Drives for Storage System. No Disk Drives found for this Storage System: *value*.

Explanation

The disk drives for this storage system cannot be discovered. This subsystem's SMI-S 1.0.2 CIM Provider is likely not implementing the disk drive subprofile.

Action

Contact the storage system vendor and request any SMI-S 1.0.2 CIM Provider updates which may include updates to the disk drive Subprofile of the SMI-S 1.0.2 CIM Provider.

BTM0557W Cannot get Disk Drives for this Storage System: *value*.

Explanation

An error occurred while retrieving disk drives associated with this storage system.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0558I Number of Volumes Found Currently: *value*. Continuing to Probe Volumes.

Explanation

The Probe is finding the Volumes. This status update is to inform how many volumes have been processed at this point during the probe.

BTM0559I Probing Volumes for StoragePool: *value*.

Explanation

The probe is finding the volumes for this storage pool.

BTM0560I Probing Volumes for Storage System: *value*.

Explanation

The probe is finding the volumes for this storage system.

BTM0561I Probing Disks for Storage System: *value*.

Explanation

The probe is finding the disks for this storage system.

BTM0562I Probing Storage Pools for Storage System: *value*.

Explanation

The probe is finding the storage pools for this storage system.

BTM0563I Probing properties of Storage System: *value*.

Explanation

The probe is finding the properties for this storage system.

BTM0564W More than one CIM_Product indirectly associated to the following Device: *value*.

Explanation

More than one CIM_Product is indirectly associated with the device. Inaccurate manufacturer, model number, serial number, or firmware might be displayed for this device.

Action

Contact the storage system vendor and request any SMI-S 1.0.2 CIM Provider updates that might include updates to the PhysicalPackage component of the CIM Provider for the device.

BTM0565W Exception caught while getting CIM_Product info for Storage System: *value*.

Explanation

An error occurred while retrieving the CIM_Product info for a Storage System. The model number, serial number, or firmware might not be known for this device.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0566W Unable to get CIM_Product info for Device: *value*.

Explanation

An error occurred while retrieving the CIM_Product info for a storage system. The manufacturer, model number, serial number, or firmware might not be known for this device.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0567W Exception caught while getting CIM_Product info for Device: *value*.

Explanation

An error occurred while retrieving the CIM_Product info for a device. The manufacturer, model number, serial number, or firmware might not be known for this device.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0568I *value* Volumes Found.

Explanation

This status message is to inform you of the total number of volumes found for this storage system or storage pool.

BTM0569I *value* Disks Found.

Explanation

This status message to inform you of the total number of disks found for this storage system or storage pool.

BTM0571W Exception caught while getting Host Initiators that can access this volume: *value*.

Explanation

An error occurred while retrieving the hosts that can access a volume. It is assumed that this value is not surfaced to any hosts.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0572W Exception caught while trying to determine RAID Level for StoragePool: *value*.

Explanation

An error occurred while trying to determine the RAID Level for a Storage Pool. This Storage Pool's RAID Level may not be reported correctly.

Action

Contact the Storage System vendor and request any SMI-S 1.0.2 CIM Provider updates which might include updates to their CIM_StorageCapabilities instances.

BTM0573E Exception caught while formatting this Host Bus Adapter port World Wide Name: *value*.

Explanation

A Host Bus Adapter port World Wide Name was not formatted correctly. This Host Bus Adapter port World Wide Name has one or more LUNs assigned to it, but this assignment will not be reported.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0574W Capacity of Disk Drive is not available: *value*.

Explanation

The MaxMediaSize property for the CIM_DiskDrive is not available or is populated with invalid data. The capacity of this Disk Drive is recorded as '0'. The invalid Disk Drive capacity affects Reports of total Disk Drive Capacity.

Action

Contact the storage system vendor and request any SMI-S 1.0.2 CIM Provider updates that might include updates to the DiskDrive Subprofile of the CIM Provider for the subsystem.

BTM0575W Exception caught while getting Host Initiators access to Volumes through this View: *value*.

Explanation

An error occurred while determining a Host Initiators' access to Volumes. One or more Host Initiator assignments to Volumes may be missing from reports. If any of the storage system's volumes are assigned to an IBM SAN Volume Controller, the storage of those volumes may be counted twice in the System-wide -> LUNs report.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0576I Probing Views of Host Initiator access to Volumes.

Explanation

The probe is finding the Host Initiator access to Volumes.

BTM0577I *value* Views Found.

Explanation

This status message to inform you of the total number of Views for Host Initiator access to Volumes that are found for this storage system.

BTM0578E Unable to connect to SMI-S provider. None of the default namespaces are valid for this SMI-S provider.

Explanation

None of the namespaces in namespace.config are valid for this SMI-S provider.

Action

Check the documentation for this SMI-S CIM Provider, or contact the storage system vendor. Get the namespace or namespaces to be used for this SMI-S CIM Provider. Add the namespace or namespaces to the top of the list in the namespace.config file. Each line must have only one namespace. Save and close the namespace.config file, and then try to connect to this SMI-S provider again.

BTM0600E Unable to get Array Site for Disk: *value*.

Explanation

The Array Site for this disk could not be determined.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0601E Error calling extrinsic method *value rc = value*: A timeout occurred trying to call the method.

Explanation

An error occurred while trying to call the specified extrinsic method. A timeout occurred trying to call the method rc.

Action

Resubmit the file system extension request when there is less activity on the storage system. If the problem persists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0602E Error calling extrinsic method *value rc = value*: The instance of the Logical Device is invalid.

Explanation

An error occurred while trying to call the specified extrinsic method. The instance of the logical device is invalid.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0603E Error calling extrinsic method *value rc = value*: There is a conflict in the Device Number.

Explanation

An error occurred while trying to call the specified extrinsic method. There is a conflict in the device number.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0604E Error calling extrinsic method *value rc = value*: A Device Number parameter must be provided.

Explanation

An error occurred while trying to call the specified extrinsic method. A device number parameter must be provided.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0605E Error calling extrinsic method *value rc = value*: A null Device Number is required by the device.

Explanation

Error calling the specified extrinsic method. A null Device Number is required by the device.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0606E Error calling extrinsic method *value rc = value*: The device is busy.

Explanation

Error calling the specified extrinsic method. The device is busy.

Action

Resubmit the file system extension request when there is less activity on the storage system, and if the problem persists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0607E Error calling extrinsic method *value rc = value*: The Protocol Controller is invalid.

Explanation

Error calling the specified extrinsic method. The Protocol Controller is invalid.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0608E Error calling extrinsic method *value rc = value*: The volume types are invalid.

Explanation

An error occurred while trying to call the specified extrinsic method. The volume types are not valid.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0609E Error calling extrinsic method *value rc = value*: One or more parameters are in the wrong System Scope.

Explanation

An error occurred while trying to call the specified extrinsic method. One or more parameters are in the wrong system scope.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0610E Error calling extrinsic method *value rc = value*: The controller needs to be created first.

Explanation

An error occurred while trying to call the specified extrinsic method. The controller needs to be created first.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0611E Error calling extrinsic method *value rc = value*: The ESSCLI call to create the volume access failed.

Explanation

An error occurred while trying to call the specified extrinsic method. The ESSCLI call to create the volume access failed.

Action

Check the cimom.log file where the SMI-S provider was installed for any messages related to the error occurring within the same time frame. Contact IBM customer technical support with the information gathered from the cimom.log as well as this message.

Related reference

- [Getting support](#)

BTM0612E Error calling extrinsic method *value rc = value*: The ESSCLI call to list the volume access failed.

Explanation

An error occurred while trying to call the specified extrinsic method. The ESSCLI call to list the volume access failed.

Action

Check the cimom.log file where the SMI-S provider was installed for any messages related to the error occurring within the same time frame. Contact IBM customer technical support with the information gathered from the cimom.log as well as this message.

Related reference

- [Getting support](#)

BTM0613E Error calling extrinsic method *value rc = value*: The instance of the Logical Device is not associated with the Controller.

Explanation

An error occurred while trying to call the specified extrinsic method. The instance of the logical device is not associated with the controller.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0614E Error calling extrinsic method *value rc = value*: The subject is not supported.

Explanation

An error occurred while trying to call the specified extrinsic method. The subject is not supported.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0615E Error calling extrinsic method *value rc = value*: The Privilege is not supported.

Explanation

An error occurred while trying to call the specified extrinsic method. The privilege is not supported.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0616E Error calling extrinsic method *value rc = value*: The Target is not supported.

Explanation

An error occurred while trying to call the specified extrinsic method. The target is not supported.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0617E Error calling extrinsic method *value rc = value*: A null parameter is not supported.

Explanation

An error occurred while trying to call the specified extrinsic method. A null parameter is not supported.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0618E Error calling extrinsic method *value rc = value*: Configuration Service is in use.

Explanation

Error calling the specified extrinsic method. Configuration Service is in use.

Action

Resubmit the file system extension request when there is less activity on the storage system, and if the problem persists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0619E Error calling extrinsic method *value rc = value*: The size is invalid.

Explanation

Error calling the specified extrinsic method. The size is invalid.

Action

Contact IBM customer technical support.

Resubmit the file system extension request with a different file extension size that is supported by the storage subsystem. If the problem persists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0620E Error calling extrinsic method *value rc = value*: The Element Type is invalid.

Explanation

Error calling the specified extrinsic method. The Element Type is invalid.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0621E Error calling extrinsic method *value rc = value*: The Goal is invalid.

Explanation

An error occurred while trying to call the specified extrinsic method. The goal is invalid.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0622E Error calling extrinsic method *value rc = value*: The Storage Pool is invalid.

Explanation

An error occurred while trying to call the specified extrinsic method. The storage pool is invalid.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0623E Error calling extrinsic method *value rc = value*: The redundancy for the Storage Pool is invalid.

Explanation

An error occurred while trying to call the specified extrinsic method. The redundancy for the storage pool is invalid.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0624E Error calling extrinsic method *value rc = value*: The requested Data Type does not match the Data Type for the Storage Pool.

Explanation

An error occurred while trying to call the specified extrinsic method. The requested data type does not match the data type for the storage pool.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0625E Error calling extrinsic method *value rc = value*: The Data Type is invalid.

Explanation

Error calling the specified extrinsic method. The Data Type is invalid.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0626E Error calling extrinsic method *value rc = value*: The Element is invalid.

Explanation

Error calling the specified extrinsic method. The Element is invalid.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0627E Error calling extrinsic method *value rc = value*: No parameters were specified for the modification.

Explanation

Error calling the specified extrinsic method. No parameters were specified for the modification.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0628E Error calling extrinsic method *value rc = value*: Unable to create volume.

Explanation

An error occurred while trying to call the specified extrinsic method. A volume could not be created.

Action

Check the cimom.log file where the SMI-S provider was installed for any messages related to the error occurring within the same time frame. Contact IBM customer technical support with the information gathered from the cimom.log as well as this message.

Related reference

- [Getting support](#)

BTM0629E Error calling extrinsic method *value rc = value*: The LSS already contains the maximum number of volumes.

Explanation

An error occurred while trying to call the specified extrinsic method. The requested logical subsystem already contains the maximum number of volumes allowed. This problem occurs when there are enough volume addresses in the logical subsystem to satisfy the request, but the requested volume size is beyond the capacity allowed by the logical subsystem.

Action

If all the model LUNs used to provision a file system are exhausted, create a policy where the model LUNs outside the volume group are used. If that is still not sufficient, assign a model LUN to the host from a logical subsystem with sufficient space. If that does not resolve the problem, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0630E Error calling extrinsic method *value rc = value*: There are not enough volume addresses in the LSS.

Explanation

An error occurred while trying to call the specified extrinsic method. The requested number of volume addresses exceeds the maximum number of volumes allowed in the given logical subsystems. This problem occurs when there are not enough volume addresses within the logical subsystem to satisfy the request. If the volume being created will be the last available volume address within the logical subsystem, it must be the same size of the remaining free space within the logical subsystem. Otherwise, this error occurs.

Action

If all the model LUNs used to provision a file system are exhausted, create a policy where the model LUNs outside the volume group are used. If that is still not sufficient, assign a model LUN to the host from a logical subsystem with sufficient space and available volume addresses. If that does not resolve the problem, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0631E Error calling extrinsic method *value rc = value*: The Identification parameter is missing or not unique.

Explanation

An error occurred while trying to call the specified extrinsic method. The identification parameter is missing or not unique.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0632E Error calling extrinsic method *value rc = value*: A null Ports parameter is required by the Controller.

Explanation

An error occurred while trying to call the specified extrinsic method. A null Ports parameter is required by the controller.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0633E Error calling extrinsic method *value rc = value*: The Controller is busy.

Explanation

An error occurred while trying to call the specified extrinsic method. The controller is busy.

Action

Resubmit the file system extension request when there is less activity on the storage system. If the problem persists, contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0634E Error calling extrinsic method *value rc = value*: The Identity is invalid.

Explanation

An error occurred while trying to call the specified extrinsic method. The identity is invalid.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0635E Error calling extrinsic method *value rc = value*: The Element Name is invalid.

Explanation

An error occurred while trying to call the specified extrinsic method. The element name is invalid.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0636E Error calling extrinsic method *value rc = value*: The Protocol is invalid.

Explanation

An error occurred while trying to call the specified extrinsic method. The protocol is invalid.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0637E Error calling extrinsic method *value rc = value*: The Privilege is invalid.

Explanation

An error occurred while trying to call the specified extrinsic method. The privilege is invalid.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0638E Error calling extrinsic method *value rc = value*: The Ports are invalid.

Explanation

An error occurred while trying to call the specified extrinsic method. The Ports are invalid.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

BTM0639E Error calling extrinsic method *value rc = value*: The host connection could not be deleted.

Explanation

An error occurred while trying to call the specified extrinsic method. The host connection could not be deleted.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0640E Error calling extrinsic method *value rc = value*: The host connection could not be created.

Explanation

An error occurred while trying to call the specified extrinsic method. The host connection could not be created.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0641E Error calling extrinsic method *value rc = value*: The host connection could not be set.

Explanation

An error occurred while trying to call the specified extrinsic method. The host connection could not be set.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0642E Error calling extrinsic method *value rc = value*: No Ports are available in this configuration.

Explanation

An error occurred while trying to call the specified extrinsic method. No ports are available in this configuration.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0701I Probing Managed Disks for Managed Disk Group: *value*

Explanation

The probe is finding the managed disks for this managed disk group.

BTM0702I Number of Managed Disks currently found: *value*.
Continuing to probe managed disks.

Explanation

The probe is finding the managed disks. This status update is to inform you of the number of managed disks that have been processed at this point during the probe.

BTM0703I *value* Managed Disks found.

Explanation

This is the total number of managed disks found on the corresponding SAN Volume Controller cluster.

BTM0704I Probing Virtual Disks for Cluster: *value*

Explanation

The probe is finding the virtual disks for this SAN Volume Controller cluster.

BTM0705I Number of Virtual Disks currently found: *value*.
Continuing to probe Virtual Disks.

Explanation

The probe is finding the virtual disks. This status update is to inform of the number of virtual disks that have been processed at this point during the probe.

BTM0706I *value* Virtual Disks found.

Explanation

This is the total number of Virtual disks found on the corresponding SAN Volume Controller cluster.

BTM0707I Probing Virtual Disks for Managed Disk Group: *value*

Explanation

The probe is finding the virtual disks for this managed disk group.

BTM0708I Probing Managed Disks for Cluster: *value*

Explanation

The probe is finding the managed disks for this SAN Volume Controller cluster.

BTM0709I Probing Managed Disks for Backend Controller: *value*

Explanation

The probe is finding the managed disks for this backend controller that has its storage virtualized by the SAN Volume Controller cluster.

BTM0710E Unable to retrieve data for Managed Disk: *value*

Explanation

The probe was not able to retrieve data for this managed disk.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0711E Unable to retrieve data for Managed Disk Group: *value*

Explanation

The probe was not able to retrieve data for this managed disk group.

Related reference

- [Getting support](#)

BTM0712E Unable to retrieve data for Virtual Disk: *value*

Explanation

The probe was not able to retrieve data for this virtual disk.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0713E Unable to retrieve data for Backend Controller: *value*

Explanation

The probe was not able to retrieve data for this backend controller.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0714E Unable to retrieve data for FC Port: *value*

Explanation

The probe was not able to retrieve data for this FC port.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0715E Unable to retrieve data for *value* Managed Disk(s) among the *value* Managed Disks found.

Explanation

This is the number of managed disks out of the total number of managed disks found for this SAN Volume Controller cluster that the probe was not able to retrieve data for.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0716E Unable to retrieve data for *value* Virtual Disk(s) among the *value* Virtual Disks found.

Explanation

This is the number of virtual disks out of the total number of virtual disks found for this SAN Volume Controller cluster that the probe was not able to retrieve data for.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BTM0717E Unable to retrieve FC Ports for Cluster: *value*

Explanation

The probe was not able to retrieve the FC ports for this cluster.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

BWN - Disk User Interface messages

- **BWN000000E** An object must be selected.
- **BWN000200E** The minimum size of the volume cannot be less than minimum volume size.
- **BWN000201E** The volume volume name cannot be deleted because there are host ports assigned to it.
- **BWN000202E** The volume volume name cannot be deleted, because volume deletion is not supported by this storage subsystem.
- **BWN000203E** No storage pool is available.
- **BWN000204E** The maximum size of the volume cannot be greater than maximum volume size.
- **BWN000205E** The maximum size of the XIV volume cannot be greater than maximum volume size. The size of the volume must be changed.
- **BWN000206E** The number of volumes must be reduced so that volume(s) with selected size can be created.
- **BWN000207E** The volume deletion failed. Please check job log for detailed information about the error.
- **BWN000208E** Invalid characters in volume name. Volume name can only contain: A-Z, a-z, 0-9, -, ., _ and space.

- [BWN000300E A valid quantity must be selected. It must be between 1 and max quantity.](#)
- [BWN000301E A valid size must be selected. It must be between 0 and max size.](#)
- [BWN000302E The maximum virtual disk size must be less than or equal to the available capacity \(available capacity\) for one virtual disk in the managed-disk group. If multiple virtual disks are to be created, the virtual-disk size must be less than or equal to the available capacity divided by the number of virtual disks.](#)
- [BWN000303E The number of virtual disks is invalid. The maximum number of virtual disks that can be created is \(available quantity\).](#)
- [BWN000304E At least one managed disk must be selected.](#)
- [BWN000305E The virtual disk vdisk name cannot be deleted because there are host ports assigned to it.](#)
- [BWN000306E The virtual disk name is not valid.](#)
- [BWN000307E Sequential virtual disks and multiple managed disks are selected, but round-robin assignment is not specified.](#)
- [BWN000308E The number of selected mdisks must be equal to the number of vdisks to be created - vdisk no..](#)
- [BWN000309E The selected managed disks could not be added to the managed-disk group.](#)
- [BWN000310E The selected host type does not match the host type of the selected host ports to be assigned.](#)
- [BWN000311E The selected host ports have identical WWPNs. Select the host ports with different WWPNs.](#)
- [BWN000312E The maximum virtual-disk real size must be less than or equal to the available capacity \(available capacity\) in the managed-disk group. If multiple virtual disks are to be created, the virtual-disk real size must be less than or equal to the available capacity divided by the number of virtual disks.](#)
- [BWN000313E When creating Space Efficient virtual-disks, the maximum virtual-disk size must not exceed maximum size.](#)
- [BWN000314E The virtual-disk warning size must be greater than 0 and cannot exceed 100 percent.](#)
- [BWN000315W This SMI-S provider is already defined with the same parameters. Would you like to save it anyway?](#)
- [BWN000316W This IBM Spectrum Control Server is already defined with the same parameters. Would you like to save it anyway?](#)
- [BWN000317W Testing SMI-S provider connectivity can take up to several minutes in case of an incorrectly entered port number, network problems or an unpassed firewall. Would you like to continue anyway?](#)
- [BWN000318E No managed disk is found for the selected mdisk group.](#)
- [BWN000319E The length of the generated virtual disk name \(virtual disk \) exceeds the maximum permitted length \(maximum length of characters \).](#)
- [BWN000600E Some of the selected host ports do not have a host connection configured on the subsystem\and require a host type to be specified. Select the appropriate host type to be used.](#)

BWN000000E An object must be selected.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000200E The minimum size of the volume cannot be less than *minimum volume size*.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000201E The volume *volume name* cannot be deleted because there are host ports assigned to it.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000202E The volume *volume name* cannot be deleted, because volume deletion is not supported by this storage subsystem.

Explanation

ESS 2105 volumes cannot be deleted.

Action

BWN000203E No storage pool is available.

Explanation

Either no storage pool has been selected, or there is no storage pool available for the specified RAID level.

Action

Make sure that a storage pool is selected. The storage pool must have available capacity to create volumes.

BWN000204E The maximum size of the volume cannot be greater than *maximum volume size*.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000205E The maximum size of the XIV volume cannot be greater than *maximum volume size*. The size of the volume must be changed.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000206E The number of volumes must be reduced so that volume(s) with selected size can be created.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000207E The volume deletion failed. Please check job log for detailed information about the error.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000208E Invalid characters in volume name. Volume name can only contain: A-Z, a-z, 0-9, _, -, ~, . and space.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000300E A valid quantity must be selected. It must be between 1 and *max quantity*.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000301E A valid size must be selected. It must be between 0 and *max size*.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000302E The maximum virtual disk size must be less than or equal to the available capacity (*available capacity*) for one virtual disk in the managed-disk group. If multiple virtual disks are to be created, the virtual-disk size must be less than or equal to the available capacity divided by the number of virtual disks.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000303E The number of virtual disks is invalid. The maximum number of virtual disks that can be created is (*available quantity*).

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000304E At least one managed disk must be selected.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000305E The virtual disk *vdisk name* cannot be deleted because there are host ports assigned to it.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000306E The virtual disk name is not valid.

Explanation

The name must adhere to the following conventions:

- One to 15 characters in length
- Permissible characters include upper-case letters (A-Z), lower-case letters (a-z), numerals (0-9), a hyphen (-), and an underscore (_).
- The first character cannot be a numeral.
- The name cannot begin with an abbreviation commonly used to specify the object type, for example, the name cannot begin with VDisk, VDISK, or vdisk.

Action

Retype the name and try again.

BWN000307E Sequential virtual disks and multiple managed disks are selected, but round-robin assignment is not specified.

Explanation

To create sequential virtual disks without using round-robin assignment, you must select only one managed disk. To use multiple managed disks, select round-robin assignment.

Action

Reselect the virtual-disk creation properties and try again.

BWN000308E The number of selected mdisks must be equal to the number of vdisks to be created - *vdisk no.*

Explanation

To create sequential virtual disks using robin-robin assignment, the number of virtual disks must equal the number of managed disks.

Action

Reselect the virtual-disk creation properties and try again.

BWN000309E The selected managed disks could not be added to the managed-disk group.

Explanation

The selected managed disks could not be added to the managed-disk group. This might be due to one or more of the following conditions:

- The cluster that contains the managed disks is not stable.
- A managed disk already is part of a managed-disk group.
- A selected managed disk does not exist.
- Too few or too many managed disks were selected.
- A selected managed disk is in managed or image mode.
- A managed disk was destroyed.

Action

Select different managed disks and try again.

BWN000310E The selected host type does not match the host type of the selected host ports to be assigned.

Explanation

The host type of selected host ports is not same as the host type selected. Assigning volumes to selected host port will result in error.

Action

All the host ports that do not match the selected host types should be unselected by moving these back to Available Ports list.

BWN000311E The selected host ports have identical WWPNs. Select the host ports with different WWPNs.

Explanation

The selected host ports have identical WWPNs. Assigning more than 1 host port with same WWPN will result in error.

Action

Host ports with identical WWPNs should be unselected by moving these back to Available Ports list.

BWN000312E The maximum virtual-disk real size must be less than or equal to the available capacity (*available capacity*) in the managed-disk group. If multiple virtual disks are to be created, the virtual-disk real size must be less than or equal to the available capacity divided by the number of virtual disks.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000313E When creating Space Efficient virtual-disks, the maximum virtual-disk size must not exceed *maximum size*.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000314E The virtual-disk warning size must be greater than 0 and cannot exceed 100 percent.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000315W This SMI-S provider is already defined with the same parameters. Would you like to save it anyway?

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000316W This IBM Spectrum Control Server is already defined with the same parameters. Would you like to save it anyway?

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000317W Testing SMI-S provider connectivity can take up to several minutes in case of an incorrectly entered port number, network problems or an unpassed firewall. Would you like to continue anyway?

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000318E No managed disk is found for the selected mdisk group.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

BWN000319E The length of the generated virtual disk name (*virtual disk*) exceeds the maximum permitted length (*maximum length of characters*).

Explanation

For multiple virtual disk, a number is added to the end of the virtual disk name. This number increases the virtual disk name length and could exceed the maximum permitted length.

Action

Retype the name and try again.

BWN000600E Some of the selected host ports do not have a host connection configured on the subsystem\ and require a host type to be specified. Select the appropriate host type to be used.

Explanation

Since the selected host ports are not known to the subsystem yet, the subsystem requires a host type to be specified in order to create a new host connection and establish a functional communication between host and host port. With no or unapplicable host type selected, the connection is likely to be non-functional.

Action

A host type must be selected that matches the characteristics of all host ports to be assigned.

CMMNP - Command Line Interface (CLI) infrastructure messages

- [CMMNP2001I Nothing to modify.](#)
- [CMMNP2002I Unsupported VALUE_0 command completed successfully.](#)
- [CMMNP2900I Command "VALUE_0" aborted.](#)
- [CMMNP4500W No VALUE_0 instances found in the system.](#)
- [CMMNP9002E Cannot modify. VALUE_0 "VALUE_1" does not exist.](#)
- [CMMNP9003E No VALUE_0 instancesVALUE_1 found that match criteria: VALUE_2.](#)
- [CMMNP9004E VALUE_0 "VALUE_1" does not exist.](#)
- [CMMNP9005E Unsupported VALUE_0 command failed with a value VALUE_1](#)
- [CMMUI4444E User name not specified.](#)
- [CMMUI9000E \[3\]An unknown value "VALUE_0" for command "VALUE_1" was returned.](#)
- [CMMUI9001E unknown](#)
- [CMMUI9006E \[3\]Command failed to execute correctly.](#)
- [CMMUI9007E \[3>Password file access error: VALUE_0.](#)
- [CMMUI9008E \[3\]Malformed password file. First line of the file requires a colon delimited user:password string](#)
- [CMMUI9010E \[1\]Invalid command: "VALUE_0" not found.](#)
- [CMMUI9011E \[21\]Invalid flag: "VALUE_0".](#)
- [CMMUI9012E \[21\]Value "VALUE_0" for flag "-VALUE_1" is formatted incorrectly.](#)
- [CMMUI9013E \[21\]Missing parameter specifier after "-"](#)
- [CMMUI9014E \[21\]Flag "VALUE_0" already specified.](#)
- [CMMUI9015E \[21\]Flag "VALUE_0" missing required value.](#)

- [CMMUI9016E \[21\]Invalid value for VALUE_0: VALUE_1.](#)
- [CMMUI9017E \[21\]The VALUE_0 flag cannot be used when the VALUE_1 option is specified.](#)
- [CMMUI9018E \[21\]Command "VALUE_0" formatted incorrectly.](#)
- [CMMUI9019E \[21\]Missing required parameter: "VALUE_0"](#)
- [CMMUI9020E \[21\]"VALUE_0" is mutually exclusive of "VALUE_1"](#)
- [CMMUI9021E \[21\]VALUE_0 exceeds the maximum allowable value of VALUE_1 for parameter "VALUE_2"](#)
- [CMMUI9022E \[21\]VALUE_0 does not meet the minimum allowable value of VALUE_1 for parameter "VALUE_2"](#)
- [CMMUI9023E \[21\]Unmatched VALUE_0 characters](#)
- [CMMUI9024E \[21\]Invalid value for VALUE_0: exceeds VALUE_1 characters](#)
- [CMMUI9025E \[21\]Value "VALUE_0" for argument "VALUE_1" invalid](#)
- [CMMUI9026E VALUE_0 "VALUE_1" does not exist.](#)
- [CMMUI9027E \[21\]Value "VALUE_0" cannot be accepted with any other value for the "-VALUE_1" flag.](#)
- [CMMUI9028E \[3\]The help page for command "VALUE_0" does not exist.](#)
- [CMMUI9029E \[21\]It is required to specify parameter "VALUE_1" when using parameter "VALUE_0"](#)
- [CMMUI9030E File "VALUE_0" doesn't exist.](#)
- [CMMUI9031E \[21\]Parameter "VALUE_0" cannot be used in the same command as parameter "VALUE_1".](#)
- [CMMUI9032E VALUE_0 "VALUE_1" already exists.](#)
- [CMMUI9033E \[21\]Value "VALUE_0" for flag "-VALUE_1" already specified.](#)
- [CMMUI9034E \[21\]Multiple targets not allowed for command "VALUE_0"](#)
- [CMMUI9035E \[21\]You cannot specify multiple VALUE_0s when using the VALUE_1 flag.](#)
- [CMMUI9036E \[21\]Invalid value "VALUE_1" for "VALUE_0": contains unsupported characters.](#)
- [CMMUI9037E \[21\]Invalid "VALUE_0" name "VALUE_1": contains unsupported characters.](#)
- [CMMUI9038E \[21\]Invalid value for VALUE_0: value other than VALUE_1 or VALUE_2 specified.](#)
- [CMMUI9039E \[21\]Value for flag "-VALUE_0" can not contain a "VALUE_1".](#)
- [CMMUI9040E \[21\]Number of entries \(VALUE_0\) is exceeded for the "-VALUE_1" flag.](#)
- [CMMUI9041E \[21\]Entry "VALUE_0" exceeds the length limit \(VALUE_1\) for one item for the "-VALUE_2" flag.](#)
- [CMMUI9042E \[21\]Value for -VALUE_0 must be VALUE_1 the current setting of VALUE_2.](#)
- [CMMUI9043E \[21\]Unrecognized syntax error in command "VALUE_0"](#)
- [CMMUI9044E Cannot run "VALUE_0" as a command within the VALUE_1 application.Tip: Enter "help VALUE_2" for more information.](#)

CMMNP2001I Nothing to modify.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMNP2002I Unsupported *VALUE_0* command completed sucessfully.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMNP2900I Command "*VALUE_0*" aborted.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMNP4500W No *VALUE_0* instances found in the system.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMNP9002E Cannot modify. *VALUE_0* "*VALUE_1*" does not exist.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMNP9003E No *VALUE_0* instances *VALUE_1* found that match criteria: *VALUE_2*.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMNP9004E *VALUE_0* "*VALUE_1*" does not exist.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMNP9005E Unsupported *VALUE_0* command failed with a value *VALUE_1*

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI4444E User name not specified.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9000E [3]An unknown value "VALUE_0" for command "VALUE_1" was returned.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9001E unknown

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9006E [3]Command failed to execute correctly.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9007E [3]Password file access error: VALUE_0.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9008E [3]Malformed password file. First line of the file requires a colon delimited user:password string

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9010E [1]Invalid command: "VALUE_0" not found.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9011E [21]Invalid flag: "VALUE_0".

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9012E [21]Value "VALUE_0" for flag "-VALUE_1" is formatted incorrectly.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9013E [21]Missing parameter specifier after "-"

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9014E [21]Flag "VALUE_0" already specified.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9015E [21]Flag "VALUE_0" missing required value.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9016E [21]Invalid value for *VALUE_0*: *VALUE_1*.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9017E [21]The *VALUE_0* flag cannot be used when the *VALUE_1* option is specified.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9018E [21]Command "*VALUE_0*" formatted incorrectly

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9019E [21]Missing required parameter: "*VALUE_0*"

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9020E [21]"*VALUE_0*" is mutually exclusive of "*VALUE_1*"

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9021E [21]*VALUE_0* exceeds the maximum allowable value of *VALUE_1* for parameter "*VALUE_2*"

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9022E [21]VALUE_0 does not meet the minimum allowable value of VALUE_1 for parameter "VALUE_2"

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9023E [21]Unmatched VALUE_0 characters

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9024E [21]Invalid value for VALUE_0: exceeds VALUE_1 characters

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9025E [21]Value "VALUE_0" for argument "VALUE_1" invalid

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9026E VALUE_0 "VALUE_1" does not exist.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9027E [21]Value "VALUE_0" cannot be accepted with any other value for the "-VALUE_1" flag.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9028E [3]The help page for command "VALUE_0" does not exist.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9029E [21]It is required to specify parameter 'VALUE_1' when using parameter 'VALUE_0'

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9030E File 'VALUE_0' doesn't exist.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9031E [21]Parameter 'VALUE_0' cannot be used in the same command as parameter 'VALUE_1'.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9032E VALUE_0 'VALUE_1' already exists.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9033E [21]Value "VALUE_0" for flag "-VALUE_1" already specified.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9034E [21]Multiple targets not allowed for command ' 'VALUE_0' '

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9035E [21]You cannot specify multiple VALUE_0s when using the VALUE_1 flag.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9036E [21]Invalid value "VALUE_1" for "VALUE_0": contains unsupported characters.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9037E [21]Invalid "VALUE_0" name "VALUE_1": contains unsupported characters.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9038E [21]Invalid value for VALUE_0: value other than VALUE_1 or VALUE_2 specified.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9039E [21]Value for flag "-VALUE_0" can not contain a "VALUE_1".

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9040E [21]Number of entries (VALUE_0) is exceeded for the "-VALUE_1" flag.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9041E [21]Entry "VALUE_0" exceeds the length limit (VALUE_1) for one item for the "-VALUE_2" flag.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9042E [21]Value for -VALUE_0 must be VALUE_1 the current setting of VALUE_2.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9043E [21]Unrecognized syntax error in command "VALUE_0"

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9044E Cannot run ''VALUE_0'' as a command within the VALUE_1 application.Tip: Enter "help VALUE_2" for more information.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI - CIM Object Manager messages

- [CMMUI9900E User access to CIMOM server denied.](#)
- [CMMUI9901E User access to command "VALUE_0" denied.](#)
- [CMMUI9902E Invalid key in truststore.](#)
- [CMMUI9903E The IBM CIM Object Manager has encountered an internal error.](#)
- [CMMUI9904E Truststore access failure.](#)
- [CMMUI9905E Namespace not found in the CIMOM server: "VALUE_0".](#)
- [CMMUI9906E Host url unspecified to CIMOM server.](#)
- [CMMUI9907E Invalid host specified to CIMOM server: "VALUE_0".](#)
- [CMMUI9908E Could not connect to CIMOM server.](#)
- [CMMUI9909E Invalid port address for CIMOM server: "VALUE_0".](#)
- [CMMUI9910E An unexpected CIMOM based error occurred: "VALUE_0".](#)

CMMUI9900E User access to CIMOM server denied.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9901E User access to command "VALUE_0" denied.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9902E Invalid key in truststore.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9903E The IBM CIM Object Manager has encountered an internal error.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9904E Truststore access failure.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9905E Namespace not found in the CIMOM server: "VALUE_0".

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9906E Host url unspecified to CIMOM server.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9907E Invalid host specified to CIMOM server: "VALUE_0".

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9908E Could not connect to CIMOM server.

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9909E Invalid port address for CIMOM server: "VALUE_0".

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CMMUI9910E An unexpected CIMOM based error occurred: "VALUE_0".

Explanation

No additional information is available for this message.

Action

Consult the Maintenance and Problem Determination Guide for suggestions for assessing system health.

CNFG - Spectrum Control Configuration messages

- [CNFG00001E The Prefix can not be blank, or contain any of the following characters: \\:*?><|."](#)
- [CNFG00002E The prefix {0} can not be used because {1} duplicate names have been found. For example: {2}](#)
- [CNFG00003E Invalid HMC Address.](#)
- [CNFG00004E Invalid Username for subsystem {0}.](#)
- [CNFG00005E Invalid Password for subsystem {0}.](#)
- [CNFG00006E Invalid IP Address for subsystem {0}.](#)
- [CNFG00007E Invalid Admin Username for subsystem {0}.](#)
- [CNFG00008E Invalid Admin Password for subsystem {0}.](#)
- [CNFG00009E Invalid Username for subsystem {0}.](#)
- [CNFG00010E Invalid Public SSH Key for subsystem {0}.](#)
- [CNFG00011E Invalid Private SSH Key for subsystem {0}.](#)
- [CNFG00012E Invalid SMI-S provider Host.](#)
- [CNFG00013E Invalid SMI-S provider Namespace.](#)
- [CNFG00014E Invalid SMI-S provider Port Number.](#)
- [CNFG00015E Invalid Out of Band Agent Host name.](#)
- [CNFG00016W You will need to update the credentials for {0}. Do you want to update the credentials now?](#)
- [CNFG00017E Connection test failed with status: {0}](#)
- [CNFG00018E Connection test to SMI-S provider {0} failed with status: {1}. {2}](#)
- [CNFG00019E The server was unable to contact out of band fabric agent at address {0} for the following reason: {1}](#)
- [CNFG00020E The filter text \[{0}\] contains invalid characters. Click the help button for allowable characters](#)
- [CNFG00021E Only {0} log files may be opened at any one time. Please reduce the number of selected log files and try again.](#)
- [CNFG00022E Unable to retrieve job data due to an internal error. Please check Data server logs.](#)
- [CNFG00023E Invalid Public SSH Key for Storwize V7000U File Module {0}.](#)
- [CNFG00024E Invalid Username for Storwize V7000U File Module {0}.](#)
- [CNFG00025E Invalid Password for Storwize V7000U File Module {0}.](#)
- [CNFG00026E Invalid Passphrase for Storwize V7000U File Module {0}.](#)
- [CNFG00027E Connection failed for Storwize V7000U File Module. Check Management Console {0}.](#)
- [CNFG00028E Invalid Public SSH Key for IBM IBM SONAS Device {0}.](#)
- [CNFG00029E Invalid Username for IBM SONAS Device {0}.](#)

- [CNFG00030E Invalid Password for IBM SONAS Device {0}.](#)
- [CNFG00031E Invalid Passphrase for IBM SONAS Device {0}.](#)

CNFG00001E The Prefix can not be blank, or contain any of the following characters: \\/:*?><|."

Explanation

A unique prefix is required.

Action

Enter a unique prefix.

CNFG00002E The prefix {0} can not be used because {1} duplicate names have been found. For example: {2}

Explanation

A probe or an alert using the prefix has been found in the database.

Action

Enter a unique prefix.

CNFG00003E Invalid HMC Address .

Explanation

The entry for HMC Address is not valid.

Action

Enter a valid HMC Address.

CNFG00004E Invalid Username for subsystem {0}.

Explanation

The entry for Username is not valid.

Action

Enter a valid Username.

CNFG00005E Invalid Password for subsystem {0}.

Explanation

The entry for Password is not valid.

Action

Enter a valid Password.

CNFG00006E Invalid IP Address for subsystem {0}.

Explanation

The entry for IP Address is not valid.

Action

Enter a valid IP Address.

CNFG00007E Invalid Admin Username for subsystem {0}.

Explanation

The entry for Admin Username is not valid.

Action

Enter a valid Admin Username.

CNFG00008E Invalid Admin Password for subsystem {0}.

Explanation

The entry for Admin Password is not valid.

Action

Enter a valid Admin Password.

CNFG00009E Invalid Username for subsystem {0}.

Explanation

The entry for Key Username is not valid.

Action

Enter a valid Key Username.

CNFG00010E Invalid Public SSH Key for subsystem {0}.

Explanation

The entry for Public SSH Key is not valid.

Action

Enter a valid Public SSH Key.

CNFG00011E Invalid Private SSH Key for subsystem {0}.

Explanation

The entry for Private SSH Key is not valid.

Action

Enter a valid Private SSH Key.

CNFG00012E Invalid SMI-S provider Host.

Explanation

The entry for Private SMI-S provider Host is not valid.

Action

Enter a valid SMI-S provider Host.

CNFG00013E Invalid SMI-S provider Namespace.

Explanation

The entry for SMI-S provider Namespace is not valid.

Action

Enter a valid SMI-S provider Namespace.

CNFG00014E Invalid SMI-S provider Port Number.

Explanation

The entry for SMI-S provider Port Number is not valid.

Action

Enter a valid SMI-S provider Port Number.

CNFG00015E Invalid Out of Band Agent Host name.

Explanation

Invalid Out of Band Agent Host Name.

Action

Enter a valid Out of Band Agent Host name.

CNFG00016W You will need to update the credentials for {0}. Do you want to update the credentials now?

Explanation

The credentials for this subsystem need to be updated.

Action

Update credentials of the subsystem.

CNFG00017E Connection test failed with status: {0}

Explanation

The connection test performed to test the system data entered did not pass for the reason specified.

Action

Double check the values entered on the device connection panel are correct. Also verify the system is running properly and configured for the port and credentials you are specifying.

CNFG00018E Connection test to SMI-S provider {0} failed with status: {1}. {2}

Explanation

The connection test performed to test the SMI-S provider data entered did not pass for the reason specified.

Action

Double check the values entered on the SMI-S provider connection panel are correct. Also verify the SMI-S provider is running properly and configured for the port and credentials you are specifying.

CNFG00019E The server was unable to contact out of band fabric agent at address {0} for the following reason: {1}

Explanation

The connection test performed to test the agent data entered did not pass for the reason specified.

Action

Double check the values entered on the agent connection panel are correct. Also verify the device is running properly and configured for the port and credentials you are specifying.

CNFG00020E The filter text [{0}] contains invalid characters. Click the help button for allowable characters

Explanation

An invalid character was input into the filter field.

Action

Double check that the values entered in the filter field are valid. For a list of allowable characters click the help button associated with the filter field.

CNFG00021E Only {0} log files may be opened at any one time. Please reduce the number of selected log files and try again.

Explanation

The maximum number of log files that can be opened at a time has been exceeded.

Action

Please reduce the number of selected log files and try again.

CNFG00022E Unable to retrieve job data due to an internal error. Please check Data server logs.

Explanation

An unexpected error occurred when attempting to retrieve job data from the server.

Action

Refer to the Data server logs and contact Customer Support

Related reference

- [Getting support](#)
- [Default locations of log files](#)

CNFG00023E Invalid Public SSH Key for Storwize V7000U File Module {0}.

Explanation

The entry for Public SSH Key for Storwize V7000U File Module is not valid.

Action

Enter a valid Public SSH Key.

CNFG00024E Invalid Username for Storwize V7000U File Module {0}.

Explanation

The entry for Storwize V7000U File Module Username is not valid.

Action

Enter a valid Key Username for Storwize V7000U File Module.

CNFG00025E Invalid Password for Storwize V7000U File Module {0}.

Explanation

The entry for Storwize V7000U File Module Password is not valid.

Action

Enter a valid Password for Storwize V7000U File Module.

CNFG00026E Invalid Passphrase for Storwize V7000U File Module {0}.

Explanation

The entry for Storwize V7000U File Module Passphrase is not valid.

Action

Enter a valid Passphrase that matches the SSH certificate for Storwize V7000U File Module.

CNFG00027E Connection failed for Storwize V7000U File Module. Check Management Console {0}.

Explanation

Connection failed for the Storwize V7000U Management Console.

Action

Storwize V7000U File Module may not be configured. Use Management Console to check the status of the Storwize V7000U File Module.

CNFG00028E Invalid Public SSH Key for IBM IBM SONAS Device {0}.

Explanation

The entry for Public SSH Key for IBM IBM SONAS Device is not valid.

Action

Enter a valid Public SSH Key.

CNFG00029E Invalid Username for IBM SONAS Device {0}.

Explanation

The entry for IBM SONAS Device Username is not valid.

Action

Enter a valid Key Username for IBM SONAS Device.

CNFG00030E Invalid Password for IBM SONAS Device {0}.

Explanation

The entry for IBM SONAS Device Password is not valid.

Action

Enter a valid Password for IBM SoNAS Device.

CNFG00031E Invalid Passphrase for IBM SONAS Device {0}.

Explanation

The entry for IBM SONAS Device Passphrase is not valid.

Action

Enter a valid Passphrase that matches the SSH certificate for IBM SONAS Device.

DIS - Discovery messages

- [DIS0001E Discover: invalid command " command".](#)
- [DIS0001I Command command selected.](#)
- [DIS0002E Discover: invalid option " option".](#)
- [DIS0003E Discover: value missing for option " option".](#)
- [DIS0004E Command command is missing a required parameter.](#)
- [DIS0005E NetWare functionality implemented only on Windows](#)
- [DIS0006E Error occurred while processing datafile datafile.](#)
- [DIS0007E Unable to send Discovery results to the server.](#)
- [DIS0008E NAS Server server name was not found as a referenced computer in the mnttab of computer computer.](#)

DIS0001E Discover: invalid command " command".

Explanation

Invalid command specified.

Action

You must enter a valid command. Contact IBM customer technical support if you are not sure about available commands.

Related reference

- [Getting support](#)

DIS0001I *Command* command selected.

Explanation

Discovery is executing the specified command.

Action

None.

DIS0002E Discover: invalid option " *option*".

Explanation

Invalid option specified for command.

Action

You must enter a valid option. Contact IBM customer technical support if you are not sure about available command options.

Related reference

- [Getting support](#)

DIS0003E Discover: value missing for option " *option*".

Explanation

The specified option is missing a required value.

Action

You must enter a valid option value. Contact IBM customer technical support if you are not sure about available option values.

Related reference

- [Getting support](#)

DIS0004E *Command* command is missing a required parameter.

Explanation

The specified command is missing a required parameter.

Action

You must enter valid parameters for the comand. Contact IBM customer technical support if you are not sure about available parameters.

Related reference

- [Getting support](#)

DIS0005E NetWare functionality implemented only on Windows

Explanation

NetWare functionality implemented only on Windows

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

DIS0006E Error occurred while processing datafile *datafile*.

Explanation

An error occurred while processing data returned from Discovery.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

DIS0007E Unable to send Discovery results to the server.

Explanation

The Agent was unable to send Discovery results to the Data Server.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

DIS0008E NAS Server *server name* was not found as a referenced computer in the mnttab of computer *computer*.

Explanation

The NAS server entered was not found in the mnttab of the selected host.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

EMSG - DS8000 management console messages

- [EMSG0001E The DS8000 Element Managers view is not accessible because the Device Server is down.](#)
- [EMSG0002E The embedded browser widget was unexpectedly destroyed. Click OK to reset display.](#)
- [EMSG0003I The element manager is about to be removed. Once removed the element manager will not be accessible from IBM Spectrum Control. Do you wish to continue?](#)
- [EMSG0004I Connection test to the element manager element manager passed.](#)
- [EMSG0005E Connection test to the element manager element manager failed.](#)
- [EMSG0006I SMI-S provider connection was removed from element manager.](#)
- [EMSG0007E A problem occurred removing the SMI-S provider from the element manager.](#)
- [EMSG0008I The SMI-S provider has been added successfully. IBM Spectrum Control has started the discovery job for the device managed by this SMI-S provider. To check the status of the jobs, go to the IBM Spectrum Control perspective and check the following navigation tree nodes: Administrative Services -> Discovery -> CIMOM](#)
- [EMSG0009E The element manager already exists.](#)
- [EMSG0010E A problem occurred adding the element manager to IBM Spectrum Control.](#)
- [EMSG0011E Connection test to Element Manager failed.](#)
- [EMSG0012E Unable to establish an https connection to the element manager.](#)
- [EMSG0013E SMI-S provider connection was added to the element manager, however the SMI-S provider discovery job failed to launch.](#)
- [EMSG0014E An element manager named manager.name already exists.](#)
- [EMSG0015E A problem occurred updating the element manager to IBM Spectrum Control.](#)
- [EMSG0016E Connection test to SMI-S provider FAILED due to status code.](#)
- [EMSG0017E The element manager's URL is not in the correct format.](#)
- [EMSG0018E A problem occurred locating the element manager in IBM Spectrum Control.](#)
- [EMSG0019E A problem occurred locating the SMI-S provider associated with the element manager.](#)
- [EMSG0020E Unable to reset DS8000 Element Manager password.](#)
- [EMSG0021E To be added to IBM Spectrum Control, all DS8000 Element Manager software is prior to release 4.2 need Username and password.](#)
- [EMSG0022E User credentials provided for DS8000 Element Manager are not valid. Provide valid credentials.](#)
- [EMSG0023E Unable to Add or Modify DS8000 Element manager for unknown reasons.](#)
- [EMSG0024E Unable to Add or Modify DS8000 Element manager. Provide right password.](#)
- [EMSG0025E Unable to Add or Modify DS8000 Element manager. Account does not exist.](#)
- [EMSG0026E Unable to Add or Modify DS8000 Element manager. User Account is locked.](#)
- [EMSG0027E Unable to Add/Modify DS8000 Element manager. Storage Authentication Service \(SAS\) database cannot be accessed.](#)
- [EMSG0028E Unable to Add or Modify DS8000 Element manager. Parameters passed are not valid.](#)
- [EMSG0029E Unable to Add or Modify DS8000 Element manager. Storage Authentication Service\(SAS\) database login task has failed. Username or password is Incorrect.](#)
- [EMSG0030E Unable to Add or Modify DS8000 Element manager. The external user account or user group is not mapped to a DS series user role.](#)
- [EMSG0031E Unable to Add or Modify DS8000 Element manager. Token submitted to Authentication policy is not supported.](#)
- [EMSG0032E Unable to Add or Modify DS8000 Element manager. Token submitted for authentication has expired. Re-authenticate to continue.](#)
- [EMSG0033E Unable to Add or Modify DS8000 Element manager. No URL provided for Storage Authentication Service \(SAS\) policy.](#)
- [EMSG0034E Unable to Add or Modify DS8000 Element manager. The host specified in the URL is not a known host or is not reachable.](#)
- [EMSG0035E Unable to Add or Modify DS8000 Element manager. The specified truststore does not have a valid certificate for the Storage Authentication Service\(SAS\).](#)
- [EMSG0036E Unable to Add or Modify DS8000 Element manager. Connection request to the Storage Authentication Service \(SAS\) is refused.](#)
- [EMSG0037E Unable to Add or Modify DS8000 Element manager. Connection request to the Storage Authentication Service \(SAS\) has failed due to socket timeout.](#)
- [EMSG0038E Unable to Add or Modify DS8000 Element manager. Failed to login to DS8000 Element manager. Please enter valid Username and password.](#)
- [EMSG0039E Unable to Add or Modify DS8000 Element manager. The type of token submitted is not supported by the Storage Authentication Service.](#)
- [EMSG0040I The SMI-S provider Connection is about to be removed. Do you wish to continue?](#)
- [EMSG0041E The DS8000 Network server is unavailable.](#)
- [EMSG0042E You are not authorized to perform this action through the management console.](#)
- [EMSG0043E You cannot add an element manager for this DS8K as its software version is too low.](#)
- [EMSG0044E You cannot update the element manager for this DS8K as its software version is too low.](#)

EMSG0001E The DS8000 Element Managers view is not accessible because the Device Server is down.

Explanation

The DS8000 Element Managers view is hosted on the Device Server.

Action

EMSG0002E The embedded browser widget was unexpectedly destroyed. Click OK to reset display.

Explanation

The panel needs to be reset because the widget containing the element manager GUI was destroyed.

Action

EMSG0003I The element manager is about to be removed. Once removed the element manager will not be accessible from IBM Spectrum Control. Do you wish to continue?

Explanation

Removes the reference to the element manager from the IBM Spectrum Control. Once the deletion is completed the element manager will not be accessible from IBM Spectrum Control. If access to the element manager is required in future then the element manager needs to be readded.

Action

Click cancel to take no further action and maintain access to the element manager. Click OK to proceed with the removal of the element manager.

EMSG0004I Connection test to the element manager *element manager* passed.

Explanation

A connection test was successfully performed on the specified element manager with the displayed result.

Action

The connection test passed and no further action is required.

EMSG0005E Connection test to the element manager *element manager* failed.

Explanation

A connection could not be established to the specified element manager.

Action

There could be a number of reasons for the failure, including, but not restricted to the element manager's machine being down, network problems or the element manager settings in IBM Spectrum Control.

EMSG0006I SMI-S provider connection was removed from element manager.

Explanation

The SMI-S provider connection was successfully removed.

Action

No action necessary.

EMSG0007E A problem occurred removing the SMI-S provider from the element manager.

Explanation

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

EMSG0008I The SMI-S provider has been added successfully. IBM Spectrum Control has started the discovery job for the device managed by this SMI-S provider. To check the status of the jobs, go to the IBM Spectrum Control perspective and check the following navigation tree nodes: Administrative Services -> Discovery -> CIMOM

Explanation

Action

This is an internal error. Contact IBM support.

EMSG0009E The element manager already exists.

Explanation

The user is attempting to add a duplicate of an element manager already defined in IBM Spectrum Control.

Action

Change the settings so that the element manager you are trying to add is unique.

EMSG0010E A problem occurred adding the element manager to IBM Spectrum Control.

Explanation

Some unexpected error occurred when attempting to add the element manager.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

EMSG0011E Connection test to Element Manager failed.

Explanation

Either the element manager information was entered incorrectly or there was a problem communicating with the element manager.

Action

Check element manager details. Ensure that the element manager is online and that there is a working network connection between both boxes.

EMSG0012E Unable to establish an https connection to the element manager.

Explanation

The servlet received a request containing an undefined action.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0013E SMI-S provider connection was added to the element manager, however the SMI-S provider discovery job failed to launch.

Explanation

Some unexpected error occurred when the backend configuration processing a request.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0014E An element manager named *manager.name* already exists.

Explanation

The name of an element manager must be unique across the element managers defined in IBM Spectrum Control. The user has attempted to name an element manager with a name that has already been used.

Action

Choose a new name for the element manager.

EMSG0015E A problem occurred updating the element manager to IBM Spectrum Control.

Explanation

Some unexpected error occurred when attempting to update the element manager.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0016E Connection test to SMI-S provider FAILED due to *status code*.

Explanation

The servlet received a request containing an undefined action.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0017E The element manager's URL is not in the correct format.

Explanation

The element manager could not be added to the management console because the information used to construct the URL of the element manager's administration console was wrong.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0018E A problem occurred locating the element manager in IBM Spectrum Control.

Explanation

Element manager not found in the IBM Spectrum Control Database.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0019E A problem occurred locating the SMI-S provider associated with the element manager.

Explanation

Some unexpected error occurred when attempting find the element manager's SMI-S provider. It could not be found.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0020E Unable to reset DS8000 Element Manager password.

Explanation

Unable to reset DS8000 Element Manager password. Please try with right credentials.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0021E To be added to IBM Spectrum Control, all DS8000 Element Manager software is prior to release 4.2 need Username and password.

Explanation

All DS8000 Element Managers prior to release 4.2 need Username or password to be added to IBM Spectrum Control.

Action

Enter a username and password

EMSG0022E User credentials provided for DS8000 Element Manager are not valid. Provide valid credentials.

Explanation

Username or Password provided for DS8000 Element manager are not valid. Provide valid Username and password.

Action

Enter a valid username and password

EMSG0023E Unable to Add or Modify DS8000 Element manager for unknown reasons.

Explanation

Unable to Add or Modify DS8000 Element manager for Unknown reasons.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0024E Unable to Add or Modify DS8000 Element manager. Provide right password.

Explanation

Provide correct password credential to Add or Modify DS8000 Element Manager.

Action

Provide a valid password

EMSG0025E Unable to Add or Modify DS8000 Element manager. Account does not exist.

Explanation

Provided account does not exist at DS8000 Element Manager.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0026E Unable to Add or Modify DS8000 Element manager. User Account is locked.

Explanation

Provided account is locked at DS8000 Element Manager.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0027E Unable to Add/Modify DS8000 Element manager. Storage Authentication Service (SAS) database cannot be accessed.

Explanation

Database is not accessible at DS8000 Element Manager.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0028E Unable to Add or Modify DS8000 Element manager. Parameters passed are not valid.

Explanation

Parameters recieved by DS8000 Element Manager are not valid.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0029E Unable to Add or Modify DS8000 Element manager. Storage Authentication Service(SAS) database login task has failed.
Username or password is Incorrect.

Explanation

DS8000 Element manager unable to log in to Database.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

EMSG0030E Unable to Add or Modify DS8000 Element manager. The external user account or user group is not mapped to a DS series user role.

Explanation

User not mapped to an ESSNI group at DS8000 Element Manager.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

EMSG0031E Unable to Add or Modify DS8000 Element manager. Token submitted to Authentication policy is not supported.

Explanation

Token passed to DS8000 Element Manager is not supported by its Authentication policy.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

EMSG0032E Unable to Add or Modify DS8000 Element manager. Token submitted for authentication has expired. Re-authenticate to continue.

Explanation

Unable to Add or Modify DS8000 Element manager. Token submitted for authentication has expired. Re-authenticate to continue.

Action

Re-authenticate to continue.

EMSG0033E Unable to Add or Modify DS8000 Element manager. No URL provided for Storage Authentication Service (SAS) policy.

Explanation

No Remote Location for Storage Authentication Service(SAS) policy provided.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

EMSG0034E Unable to Add or Modify DS8000 Element manager. The host specified in the URL is not a known host or is not reachable.

Explanation

Unreachable or Unknown URL specified for DS8000 Element manager.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

EMSG0035E Unable to Add or Modify DS8000 Element manager. The specified truststore does not have a valid certificate for the Storage Authentication Service (SAS).

Explanation

The specified truststore does not have a valid certificate for the Storage Authentication Service (SAS).

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

EMSG0036E Unable to Add or Modify DS8000 Element manager. Connection request to the Storage Authentication Service (SAS) is refused.

Explanation

A Connection request from DS8000 Element Manager to the Storage Authentication Service (SAS) is refused.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0037E Unable to Add or Modify DS8000 Element manager. Connection request to the Storage Authentication Service (SAS) has failed due to socket timeout.

Explanation

Connection request to the Storage Authentication Service (SAS) has failed due to a socket timeout at DS8000 Element manager.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0038E Unable to Add or Modify DS8000 Element manager. Failed to login to DS8000 Element manager. Please enter valid Username and password.

Explanation

Failed to login to DS8000 Element manager. Enter a valid username and password

Action

Enter a valid username and password to continue.

EMSG0039E Unable to Add or Modify DS8000 Element manager. The type of token submitted is not supported by the Storage Authentication Service.

Explanation

The token type submitted is not compatible with the configuration of the ESS Client on the DS8000 Element manager.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

EMSG0040I The SMI-S provider Connection is about to be removed. Do you wish to continue?

Explanation

Removes SMI-S provider Connection to the element manager from the IBM Spectrum Control.

Action

Click cancel to take no further action and retain SMI-S provider Connection. Click OK to proceed with the removal of the SMI-S provider Connection.

MSG0041E The DS8000 Network server is unavailable.

Explanation

The connection to the DS8000 Network server is established immediately after the user login to the GUI application. The DS8000 Network server cannot be reached over the TCP/IP network or over a local connection. The DS8000 Network server might not be enabled, a network problem might exist, or the DS8000 Network server might be refusing connections because too many clients are connected.

Action

Ensure that the network is functioning properly and the DS8000 network server is enabled. If too many clients are connected, wait until the DS8000 Network server has fewer clients. Resubmit the task. If the problem persists, log out of the GUI application and log in again and resubmit the task.

MSG0042E You are not authorized to perform this action through the management console.

Explanation

You are attempting to perform an action for which you have no authorization.

Action

Contact your system administrator to change your authorization level. Resubmit the task.

MSG0043E You cannot add an element manager for this DS8K as its software version is too low.

Explanation

The DS8K must be release 3 or later to add using Add Element Manager.

Action

Upgrade the DS8K to a more software version greater than R5.3.

MSG0044E You cannot update the element manager for this DS8K as its software version is too low.

Explanation

The DS8K must be release 3 or later to update using Update Element Manager.

Action

Upgrade the DS8K to a more software version greater than R5.3.

GEN - General Spectrum Control messages

- [GEN0001E command name\(command arguments\) FAILED.](#)
- [GEN0002E Unable to create temporary file file name.](#)
- [GEN0003E Unable to lock temporary file file name.](#)
- [GEN0004E accept\(\) fails -- error message.](#)
- [GEN0005E Cannot get input stream from host <computer name>.](#)
- [GEN0006E Cannot get output stream to host <computer name>.](#)
- [GEN0007E Cannot deserialize from host <computer name>.](#)
- [GEN0008E Cannot read from host <computer name>.](#)
- [GEN0009E Cannot serialize to host <computer name>.](#)

- [GEN0010E Cannot write to host <computer name>.](#)
- [GEN0011E Cannot open object stream from host <computer name>.](#)
- [GEN0012E Socket to host <computer name> closed prematurely.](#)
- [GEN0013E Object read from host <computer name> is not Request -- it is.](#)
- [GEN0014E Object read from host <computer name> is not Response -- it is.](#)
- [GEN0015E Socket to host <computer name> -- class not found.](#)
- [GEN0016I Above error occurred sending Request \(request type, request subtype\).](#)
- [GEN0017I Above error occurred sending Request \(request type, request subtype\).\(phase value\).](#)
- [GEN0018I Response received from host <computer name>.](#)
- [GEN0019I Above error occurred responding to Request \(request type, request subtype\).](#)
- [GEN0020I Above error occurred reading data for Request \(request type, request subtype\).](#)
- [GEN0021E Error closing socket to host <computer name>.](#)
- [GEN0023E Cannot rename old file name to new file name.](#)
- [GEN0024I On socket to host <computer name>.](#)
- [GEN0025E Request\(request type, request subtype\) was incomplete, but response said complete.](#)
- [GEN0026E Request\(request type, request subtype\) was complete, but response said incomplete.](#)
- [GEN0027E Cannot send incomplete response to single-phase Request\(request type, request subtype\).](#)
- [GEN0028E Tried to read data of single-phase Request\(request type, request subtype\).](#)
- [GEN0029E Cannot open.](#)
- [GEN0030E Cannot create listener on port port number.](#)
- [GEN0031E Unroutable type-code.](#)
- [GEN0032E Unroutable sub-type.](#)
- [GEN0033W SoTimeout failed -- host <computer name>.](#)
- [GEN0035E Cannot seek to position on file name.](#)
- [GEN0036E Cannot create.](#)
- [GEN0037E Cannot create file name.](#)
- [GEN0038W Cannot delete file name.](#)
- [GEN0039E Unknown host <computer name>.](#)
- [GEN0040E Cannot connect to <computer name>:port number>.](#)
- [GEN0041E System property <property name> is not defined.](#)
- [GEN0042E Error reading log-file file name.](#)
- [GEN0043E Cannot find SM/DMI header in low memory.](#)
- [GEN0044E Cannot seek to physical memory address address.](#)
- [GEN0045E No type-1 structure found.](#)
- [GEN0046E Obsolete type-1 structure \(no UUID\).](#)
- [GEN0047E Invalid or unset UUID.](#)
- [GEN0048E Cannot obtain system manufacturer.](#)
- [GEN0049E Cannot open log file file name.](#)
- [GEN0050E Unable to connect to database repository.](#)
- [GEN0052E Cannot open directory directory name.](#)
- [GEN0054E Error creating pipe.](#)
- [GEN0055E Error sussing pipe.](#)
- [GEN0056E Cannot perform operation <operation> on physical memory.](#)
- [GEN0057W Windows message message.](#)
- [GEN0058W Windows message message: description.](#)
- [GEN0067W Cannot write to pipe name.](#)
- [GEN0068E Cannot find binary module module name.](#)
- [GEN0069E Cannot find entry point entry point name.](#)
- [GEN0070E Cannot find privilege privilege name.](#)
- [GEN0071E Cannot open own token.](#)
- [GEN0072E Cannot assert privilege privilege name.](#)
- [GEN0073E Cannot look up network interfaces.](#)
- [GEN0074E No Ethernet cards found.](#)
- [GEN0075E GetTokenInformation\(\) failed.](#)
- [GEN0076E Not super-user.](#)
- [GEN0077E SetHandleInformation\(\) failed.](#)
- [GEN0078I Trying token-ring.](#)
- [GEN0079E socket\(\) failed.](#)
- [GEN0080E Error looking up Ethernet card.](#)
- [GEN0081W No token-ring cards found.](#)
- [GEN0082W Error looking up token-ring card.](#)
- [GEN0083I Trying Ethernet.](#)
- [GEN0084I Hardware-ID obtained.](#)
- [GEN0096I PID = process identifier.](#)
- [GEN0097E Unable to retrieve hardware-ID.](#)
- [GEN0098E Error processing request from host <computer name>, user <user name>, for service <service name>, request\(request type, request subtype\).](#)
- [GEN0099W Warning processing request from host <computer name>, user <user name>, for service <service>, request\(request type, request subtype\).](#)
- [GEN0100E Error processing request from host <computer name>, for service <service>, request\(request type, request subtype\).](#)
- [GEN0101W Warning processing request from host <computer name>, for service <service>, request\(request type, request subtype\).](#)
- [GEN0102E Missing t_ identifiers row. ID type id number.](#)
- [GEN0104E Not enough virtual memory.](#)
- [GEN0105E Unable to send internal job results to T-Time Schedule : scheduler, Job: job name Run Number: run number.](#)
- [GEN0106E Unable to obtain local hostname. hostname.](#)
- [GEN0107E Unable to obtain SNMP datagram socket.](#)
- [GEN0108E Unable to send SNMP trap datagram.](#)
- [GEN0109E Unable to deliver an e-mail to one or more of the following recipients: recipients](#)
- [GEN0110E Unable to read reply from SMTP server server name.](#)
- [GEN0111E Unable to connect to SMTP server server Authentication failed](#)
- [GEN0112E Unable to write to SMTP server server name.](#)
- [GEN0113E Unable to connect to SMTP server server name Unknown host.](#)

- [GEN0114E Unable to connect to SMTP server server name, port port number.](#)
- [GEN0115E SNMP server server name is unknown.](#)
- [GEN0125E Requested Report invalid: report name/report number.](#)
- [GEN0126W Unable to determine home directory.](#)
- [GEN0127W job creator.job name job messages will be logged to the Data Server log file.](#)
- [GEN0128E License key must be in format xxxxx-xxxxx-xxxxx-xxxxx.](#)
- [GEN0129E License key characters must be 0-9 or A-Z excluding E, I, O, U.](#)
- [GEN0130E License key checksum is incorrect - key is invalid.](#)
- [GEN0131E Unable to open file file name.](#)
- [GEN0132E Error parsing file name. Unrecognized section name: section name.](#)
- [GEN0133E Error reading file file name.](#)
- [GEN0134E Error parsing file file name Unrecognized token token name in section section name.](#)
- [GEN0135E Error - file file name appears to be truncated.](#)
- [GEN0136E Not enough agent licenses to run job.](#)
- [GEN0137E product name license has expired.](#)
- [GEN0139E product name is not installed on agent agent name.](#)
- [GEN0140E product name is not licensed for agent agent name.](#)
- [GEN0141E License key key is invalid.](#)
- [GEN0143E key license key is expired.](#)
- [GEN0148E Unknown product on computer name.](#)
- [GEN0149E product name is at release release level on agent name and at release release level on the server.](#)
- [GEN0151I product name vversion.release.modification.](#)
- [GEN0152E Error writing file name.](#)
- [GEN0153E Error serializing to file namefile name.](#)
- [GEN0154E Unable to retrieve cached report data Error reading file name.](#)
- [GEN0155E Unable to retrieve cached report data Error deserializing from file name.](#)
- [GEN0156E Unable to retrieve cached report data Premature end of file -- file name.](#)
- [GEN0157E Unable to retrieve cached report data Class class name not found restoring from file name.](#)
- [GEN0158E Unable to retrieve cached report data Object restored from class name is not file name; it's.](#)
- [GEN0159E Unable to retrieve cached report data Your request was interrupted. Processing terminated.](#)
- [GEN0160E Cached report data is no longer available on the server. Generate the report again. Cannot open.](#)
- [GEN0161E Error setting permissions on file name.](#)
- [GEN0162E Error retrieving FD flags.](#)
- [GEN0163E Error setting FD flags.](#)
- [GEN0164E Write timed out.](#)
- [GEN0165E poll\(\) failed.](#)
- [GEN0166E write\(\) failed.](#)
- [GEN0167E SO_SNDTIMEO failed.](#)
- [GEN0172E Error processing Request\(request type, request subtype\).](#)
- [GEN0173E Java Error in readObject\(\): error message.](#)
- [GEN0174E requestData is: class name.](#)
- [GEN0175E writeObject\(\) failed writing: error message.](#)
- [GEN0176E responseData was: class name.](#)
- [GEN0177E Java Error in writeObject: error message.](#)
- [GEN0178E GuiReportReq report\(report type, report subtype\).](#)
- [GEN0179E GuiListReq listRequested\(type\).](#)
- [GEN0180W A license key exists for product name, but the software is not installed.](#)
- [GEN0181E Error - duplicate rows found for agent agent name.](#)
- [GEN0182I TsName: name, Manufacturer: manufacturer, HwID: hardware id.](#)
- [GEN0184E License key is invalid for this software release.](#)
- [GEN0197E Bad magic number.](#)
- [GEN0198I product component starting.](#)
- [GEN0199E The license key for this product has been deleted.](#)
- [GEN0200E Not enough licenses to license all requested agents.](#)
- [GEN0201W operating system : License not present or expired.](#)
- [GEN0222W Cannot find binary module module name.](#)
- [GEN0223W Cannot find entry point entry point name.](#)
- [GEN0304E Failed to read the stream header.](#)
- [GEN0305E Failed to decrypt input stream.](#)
- [GEN0306E Authentication failed from host:request type request subtype.](#)
- [GEN0307E Request not allowed from a cryptable stream request type,request subtype.](#)
- [GEN0308E Failed to write the stream header.](#)
- [GEN0309E Failed to re-authenticate.](#)
- [GEN0310E Failed to encrypt input stream.](#)
- [GEN0311E Authentication failed.](#)
- [GEN0329E Failed to authenticate the user <user ID>.](#)
- [GEN0332E Unable to connect to host hostname, port port.](#)
- [GEN0335E Unable to send a test email message to server name.](#)
- [GEN0400I Probe aborted.](#)
- [GEN0400E Probe completed with errors.](#)
- [GEN0401I Probe completed successfully.](#)
- [GEN0402W Probe completed with warnings.](#)
- [GEN0403E Unable to retrieve probe definition.](#)
- [GEN0404I Probe started.](#)
- [GEN0405E No computers or storage subsystems to probe.](#)
- [GEN0406E A probe of the computer or hypervisor computer name is already in progress.](#)
- [GEN0407E Probe of storage subsystem subsystem name \(subsystem alias\) already in progress.](#)
- [GEN0408W Probe of fabric fabric WWN already in progress.](#)
- [GEN0409E Batch Report job report name already in progress.](#)
- [GEN0410E Probe of switch switch WWN already in progress.](#)

- [GEN0557E An OS error occurred.](#)
- [GEN0558E An IO error occurred.](#)
- [GEN1030E product name is not licensed on computer computer name.](#)
- [GEN1034I GeneralException message text follows:error message.](#)
- [GEN1035W storage subsystem name is no longer monitored.](#)
- [GEN6013E OS Error errno: error description](#)
- [GEN6014W OS Error errno](#)
- [GEN6015E Unable to load JAVA.DLL](#)
- [GEN6016E Unable to find symbol module symbol in JAVA.DLL](#)
- [GEN6017E Executable file is null](#)
- [GEN6018E stat\(file path\) failed.](#)
- [GEN6019E chmod\(file path\) failed.](#)
- [GEN5001W Not all statistics could be saved due to unlicensed computer\(s\). See TPCD log for details.](#)
- [GEN7111E Unable to retrieve data from the repository database.](#)
- [GEN7112E Failed to retrieve e-mail report from \(report URL\).](#)
- [GEN7113E Failed to retrieve e-mail report from the web server.](#)
- [GEN7114E The user and password combination for the outgoing e-mail server is invalid.](#)
- [GEN7115E Can't send the attached "reportName" report with the ID reportId. The size of the attachment \(totalSize KB\) exceeds the maximum size for attachments that was set on your email server. Learn More link](#)
- [GEN7116E The email address that was configured for replies to alert notifications and reports, reply_to_address value, was not accepted by the email server, server.](#)

GEN0001E *command name (command arguments)* FAILED.

Explanation

The specified command FAILED.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0002E Unable to create temporary file *file name*.

Explanation

The temporary file could not be created.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0003E Unable to lock temporary file *file name*.

Explanation

The temporary file could not be locked.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0004E accept() fails -- error message.

Explanation

The socket accept() command failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0005E Cannot get input stream from host <computer name>.

Explanation

Stream input from the host cannot be received.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0006E Cannot get output stream to host <computer name>.

Explanation

Stream output from the host cannot be received.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0007E Cannot deserialize from host <computer name>.

Explanation

The message from the host cannot be deserialized.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0008E Cannot read from host <computer name>.

Explanation

The Object Input stream from the host cannot be read.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0009E Cannot serialize to host <computer name>.

Explanation

The message from the host cannot be serialized.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0010E Cannot write to host <computer name>.

Explanation

The host cannot be written to.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0011E Cannot open object stream from host <computer name>.

Explanation

The object stream from host cannot be opened.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0012E Socket to host <computer name> closed prematurely.

Explanation

The Socket to the host was closed prematurely.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0013E Object read from host <computer name> is not Request -- it is.

Explanation

The object read from host is not a request.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0014E Object read from host <computer name> is not Response -- it is.

Explanation

The object read from the host is not a response.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0015E Socket to host <computer name> -- class not found.

Explanation

The socket to host class was not found.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0016I Above error occurred sending Request (request type, request subtype) .

Explanation

The specified error occurred while sending a request.

GEN0017I Above error occurred sending Request (*request type, request subtype*) (*phase value*).

Explanation

The specified error occurred while sending a request during the specified phase.

GEN0018I Response received from host <*computer name*>.

Explanation

A response was received from the host.

GEN0019I Above error occurred responding to Request (*request type, request subtype*).

Explanation

The specified error occurred responding to a request.

GEN0020I Above error occurred reading data for Request (*request type, request subtype*).

Explanation

The specified error occurred while reading data from a request.

GEN0021E Error closing socket to host <*computer name*>.

Explanation

An error occurred closing the socket to the host.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

GEN0023E Cannot rename *old file name* to *new file name*.

Explanation

The specified file cannot be renamed.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

GEN0024I On socket to host <computer name>.

Explanation

On socket to host.

GEN0025E Request(*request type*, *request subtype*) was incomplete, but response said complete.

Explanation

The request was incomplete, but its response said it was complete.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

GEN0026E Request(*request type*, *request subtype*) was complete, but response said incomplete.

Explanation

The request was complete, but its response said it was incomplete.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

GEN0027E Cannot send incomplete response to single-phase Request(*request type*, *request subtype*).

Explanation

An incomplete response cannot be sent to single-phased Request.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

GEN0028E Tried to read data of single-phase Request(*request type*, *request subtype*).

Explanation

An error occurred trying to read data of a single-phased Request.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0029E Cannot open .

Explanation

The specified file cannot be opened.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0030E Cannot create listener on port *port number*.

Explanation

A listener on the port number cannot be created.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0031E Unroutable type-code.

Explanation

An unroutable type-code was encountered.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0032E Unroutable sub-type.

Explanation

An unroutable sub-type was encountered.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0033W SoTimeout failed -- host *<computer name>*.

Explanation

The SoTimeout failed.

GEN0035E Cannot seek to *position* on *file name*.

Explanation

The file seek to the specified location cannot be completed.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0036E Cannot create.

Explanation

The specified file cannot be created.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0037E Cannot create *file name*.

Explanation

The specified file cannot create be created.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0038W Cannot delete *file name*.

Explanation

The specified file cannot be deleted.

GEN0039E Unknown host <computer name>.

Explanation

The host is not known host.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0040E Cannot connect to <computer name:port number>.

Explanation

The specified host and port number cannot be connected to.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0041E System property <property name> is not defined.

Explanation

The System property specified is not defined.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0042E Error reading log-file *file name*.

Explanation

An error occurred while reading the specified log-file.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0043E Cannot find SM/DMI header in low memory.

Explanation

The SM/DMI header cannot be found in low memory.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0044E Cannot seek to physical memory address *address*.

Explanation

The seek to the specified physical memory address failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0045E No type-1 structure found.

Explanation

A type-1 structure was not found.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0046E Obsolete type-1 structure (no UUID) .

Explanation

An Obsolete type-1 structure (no UUID) was found.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0047E Invalid or unset UUID.

Explanation

An invalid or unset UUID was found.

Action

This message indicates a condition that will not affect server startup and can be ignored.

GEN0048E Cannot obtain system manufacturer.

Explanation

The system manufacturer cannot be obtained.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0049E Cannot open log file *file name*.

Explanation

The log file cannot be opened. The file may have been deleted due to log file retention.

Action

The data you are viewing may be stale. Please refresh to get the current display and retry your request. If that does not resolve your file access issues, contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0050E Unable to connect to database repository.

Explanation

IBM Spectrum Control cannot communicate with the database repository. The database repository contains information that was collected about storage resources.

Action

Ensure that the database repository (DB2), Data server, and Device server are up and running. Verify that the local area network is available and a firewall is not preventing network access. Check the product log files for error messages that might help determine the problem. See the product information center to view the locations of these log files. Try the action again. If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

GEN0052E Cannot open directory *directory name*.

Explanation

The specified directory cannot be opened.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0054E Error creating pipe.

Explanation

An error occurred creating a pipe.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0055E Error sussing pipe.

Explanation

An error occurred while sussing pipe.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0056E Cannot perform operation <operation> on physical memory.

Explanation

The specified operation cannot be performed on physical memory.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0057W Windows message *message*.

Explanation

The specified Windows event has occurred.

Action

Visit the Microsoft Knowledge Base for more detailed information.

GEN0058W Windows message *message: description*.

Explanation

The specified Windows event has occurred.

Action

Visit the Microsoft Knowledge Base for more detailed information.

GEN0067W Cannot write to *pipe name*.

Explanation

The specified pipe cannot be written to.

GEN0068E Cannot find binary module *module name*.

Explanation

The specified binary module cannot be found.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)
-

GEN0069E Cannot find entry point *entry point name*.

Explanation

The specified entry point into the binary module cannot be found.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)
-

GEN0070E Cannot find privilege *privilege name*.

Explanation

The specified privilege could not be found.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)
-

GEN0071E Cannot open own token.

Explanation

The process token could not be opened.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0072E Cannot assert privilege *privilege name*.

Explanation

The specified privilege could not be asserted.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0073E Cannot look up network interfaces.

Explanation

The network interfaces lookup failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0074E No Ethernet cards found.

Explanation

No Ethernet cards were found.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0075E GetTokenInformation() failed.

Explanation

The call to GetTokenInformation() failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0076E Not super-user.

Explanation

The user is not a member of the administrators group.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0077E SetHandleInformation() failed.

Explanation

The call to SetHandleInformation() failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0078I Trying token-ring.

Explanation

Trying to locate a token-ring.

GEN0079E socket() failed.

Explanation

The call to socket() failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0080E Error looking up Ethernet card.

Explanation

An error occurred while looking up the Ethernet card.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

GEN0081W No token-ring cards found.

Explanation

No token-ring cards were found.

GEN0082W Error looking up token-ring card.

Explanation

An error occurred while looking up token-ring card.

GEN0083I Trying Ethernet.

Explanation

Trying to locate an Ethernet card.

GEN0084I Hardware-ID obtained.

Explanation

A Hardware-ID was obtained by querying the Ethernet Card.

GEN0096I PID = *process identifier*.

Explanation

This is an informational message stating the Process Identification Number.

GEN0097E Unable to retrieve hardware-ID.

Explanation

The hardware-ID was unable to be retrieved.

Action

Contact IBM customer technical support.

Related reference

-  [Getting support](#)

GEN0098E Error processing request from host <computer name>, user <user name>, for service <service name>, request(request type,

request subtype).

Explanation

An error occurred while processing request from the specified host and user, for the specified service request.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0099W Warning processing request from host *<computer name>*, user *<user name>*, for service *<service>*, request(*request type*, *request subtype*).

Explanation

A warning occurred while processing request from the specified host and user, for the specified service request.

GEN0100E Error processing request from host *<computer name>*, for service *<service>*, request(*request type*, *request subtype*).

Explanation

An error occurred while processing request from the specified host and user, for the specified service request.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0101W Warning processing request from host *<computer name>*, for service *<service>*, request(*request type*, *request subtype*).

Explanation

A warning occurred while processing request from the specified host and user, for the specified service request.

GEN0102E Missing *t_identifiers* row. ID type *id number*.

Explanation

The *t_identifiers* rows with the specified ID type were not found during an update.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0104E Not enough virtual memory.

Explanation

A system call to allocate memory failed because there is not enough virtual memory.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0105E Unable to send internal job results to T-Time Schedule : *scheduler*, Job: *job name* Run Number: *run number*.

Explanation

Unable to send internal job results to the specified T-Time Schedule, Job and Run Number.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0106E Unable to obtain local hostname. *hostname*.

Explanation

The specified local hostname was not able to be obtained.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0107E Unable to obtain SNMP datagram socket.

Explanation

The SNMP datagram socket was not able to be obtained.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0108E Unable to send SNMP trap datagram.

Explanation

The SNMP trap datagram was not able to be obtained.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0109E Unable to deliver an e-mail to one or more of the following recipients: *recipients*

Explanation

Failed to send an e-mail to one or more of the specified recipients.

Action

Check the validity of the e-mail addresses.

GEN0110E Unable to read reply from SMTP server *server name*.

Explanation

A reply was unable to be read from the specified SMTP server.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0111E Unable to connect to SMTP server *server* Authentication failed

Explanation

Failed to authenticate with the SMTP server.

Action

Verify the username and password and try again.

GEN0112E Unable to write to SMTP server *server name*.

Explanation

An error occurred while trying to write to the specified SMTP server.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0113E Unable to connect to SMTP server *server name* Unknown host.

Explanation

An error occurred while trying to connect to the specified SMTP server because the host was unknown.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0114E Unable to connect to SMTP server *server name*, port *port number*.

Explanation

IBM Spectrum Control cannot connect to the specified mail server.

Action

Verify network connectivity to the specified mail server. Ensure the server is active and the port is the unsecured mail server port.

GEN0115E SNMP server *server name* is unknown.

Explanation

The specified SNMP server is unknown.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0125E Requested Report invalid: *report name/report number*.

Explanation

The requested Report is not an available type of report.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0126W Unable to determine home directory.

Explanation

The home directory was not able to be determined.

GEN0127W *job creator.job name* job messages will be logged to the Data Server log file.

Explanation

The specified job's messages will be logged to the Data Server log file.

GEN0128E License key must be in format **xxxxx-xxxxx-xxxxx-xxxxx**.

Explanation

The license key entered was not entered correctly. The license key must be in format xxxxx-xxxxx-xxxxx-xxxxx.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0129E License key characters must be 0-9 or A-Z excluding E, I, O, U.

Explanation

The license key entered was not entered correctly. License key characters must be 0-9 or A-Z excluding E, I, O, U.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0130E License key checksum is incorrect - key is invalid.

Explanation

The license key entered was not entered correctly. The license key checksum is incorrect.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0131E Unable to open file *file name*.

Explanation

The specified file could not be opened.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0132E Error parsing *file name*. Unrecognized section name: *section name*.

Explanation

An error occurred parsing the specified file. An unrecognized section name was found.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0133E Error reading file *file name*.

Explanation

An error occurred reading the specified file.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0134E Error parsing file *file name* Unrecognized token *token name* in section *section name*.

Explanation

An error occurred parsing the specified file. An unrecognized token was found in the specified section.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0135E Error - file *file name* appears to be truncated.

Explanation

An error occurred while reading the configuration file, the file appears to be truncated.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0136E Not enough agent licenses to run job.

Explanation

There are not enough agent licenses to run job.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0137E *product name* license has expired.

Explanation

The specified license has expired.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0139E *product name* is not installed on agent *agent name*.

Explanation

The specified product is not installed on the agent.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0140E *product name* is not licensed for agent *agent name*.

Explanation

The specified product is not licensed for the specified agent.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0141E License key *key* is invalid.

Explanation

The license key specified is invalid.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0143E *key* license key is expired.

Explanation

The license key for the specified product is expired.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0148E Unknown product on *computer name*.

Explanation

An unknown product was found on the specified host.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0149E *product name* is at release *release level* on *agent name* and at release *release level* on the server.

Explanation

A release mismatch exists. The specified product release differs on the server and Agent.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0151I *product name vversion.release.modification.*

Explanation

The is an informational message displaying the Version, Release and Modification of the specified product.

GEN0152E Error writing *file name.*

Explanation

An error occurred writing to the specified file.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0153E Error serializing to file name*file name.*

Explanation

An error occurred serializing to the specified file.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0154E Unable to retrieve cached report data Error reading *file name.*

Explanation

An error occurred while reading from the specified file and so the cached report was unable to be retrieved.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0155E Unable to retrieve cached report data Error deserializing from *file name.*

Explanation

An error occurred while deserializing the data from the specified file and so the cached report was unable to be retrieved.

Action

Contact IBM customer technical support.

GEN0156E Unable to retrieve cached report data Premature end of file -- *file name*.

Explanation

An error occurred while reading from the specified file and so the cached report was unable to be retrieved.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0157E Unable to retrieve cached report data Class *class name* not found restoring from *file name*.

Explanation

An error occurred while reading from the specified file and so the cached report was unable to be retrieved.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0158E Unable to retrieve cached report data Object restored from *class name* is not *file name*; it's.

Explanation

The cached report was unable to be restored from the specified file, the report that was restored is not the correct report.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0159E Unable to retrieve cached report data Your request was interrupted. Processing terminated.

Explanation

The cached report was unable to be restored it was interrupted during processing.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0160E Cached report data is no longer available on the server. Generate the report again. Cannot open.

Explanation

The cached report data is no longer available on the server. Generate the report again.

Action

If this message is received for a batch report, try increasing the batchPartitionWaitRetryCount property in the TPCD.config file. Otherwise, contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0161E Error setting permissions on *file name*.

Explanation

An error occurred while trying to set permissions on the specified file.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0162E Error retrieving FD flags.

Explanation

An error occurred while retrieving the file descriptor flags.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0163E Error setting FD flags.

Explanation

An error occurred while setting the file descriptor flags.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0164E Write timed out.

Explanation

The write timed out.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0165E poll() failed.

Explanation

The poll() command failed.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0166E write() failed.

Explanation

The write() command failed.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0167E SO_SNDTIMEO failed.

Explanation

The SO_SNDTIMEO command failed.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN0172E Error processing Request(*request type*, *request subtype*).

Explanation

An error occurred processing the specified request.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0173E Java Error in readObject(): *error message*.

Explanation

The specified error occurred in the readObject() routine.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0174E requestData is: *class name*.

Explanation

The requestData class is as specified.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0175E writeObject() failed writing: *error message*.

Explanation

The writeObject() command failed while trying to write the specified class.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0176E responseData was: *class name*.

Explanation

The requestData class is as specified.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0177E Java Error in writeObject: *error message*.

Explanation

The writeObject() command failed while trying to write the specified class.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0178E GuiReportReq report(*report type, report subtype*).

Explanation

The requested report and report subtype.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0179E GuiListReq listRequested(*type*).

Explanation

The type of GUI List Requested.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0180W A license key exists for *product name*, but the software is not installed.

Explanation

A license key exists for the specified product, but the software is not installed.

GEN0181E Error - duplicate rows found for agent *agent name*.

Explanation

Duplicate rows were found for the specified agent.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0182I *TsName: name, Manufacturer: manufacturer, HwID: hardware id.*

Explanation

This is an informational message logging the name, manufacturer and hardware id of the agent.

GEN0184E *License key is invalid for this software release.*

Explanation

The license key entered is invalid for this software release of the product.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0197E *Bad magic number.*

Explanation

Bad magic number.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0198I *product component starting.*

Explanation

An informational message specifying which part of the product has started.

GEN0199E *The license key for this product has been deleted.*

Explanation

The license key for this product has been deleted.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0200E Not enough licenses to license all requested agents.

Explanation

More licenses are needed to license all requested agents.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0201W *operating system* : License not present or expired.

Explanation

The specified Operating System does not have a license present or it is expired.

GEN0222W Cannot find binary module *module name*.

Explanation

The specified binary module cannot be found.

Action

Contact IBM customer technical support.

GEN0223W Cannot find entry point *entry point name*.

Explanation

The specified entry point into the binary module cannot be found.

Action

Contact IBM customer technical support.

GEN0304E Failed to read the stream header.

Explanation

The stream header of the Socket Message could not be read.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0305E Failed to decrypt input stream.

Explanation

The input stream message could not be decrypted.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0306E Authentication failed from *host:request type request subtype*.

Explanation

Authentication failed from host.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0307E Request not allowed from a cryptable stream *request type,request subtype*.

Explanation

Request not allowed from a cryptable stream.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0308E Failed to write the stream header.

Explanation

Failed to write the stream header.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0309E Failed to re-authenticate.

Explanation

Failed to re-authenticate.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0310E Failed to encrypt input stream.

Explanation

Failed to encrypt input stream.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0311E Authentication failed.

Explanation

Authentication failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0329E Failed to authenticate the user <user ID>.

Explanation

An error occurred when authenticating the user to the server.

Action

Verify the user ID, password, and server specified in the Sign On dialog are correct. If the server has multiple host names, please enter the primary host name.

GEN0332E Unable to connect to host *hostname*, port *port*.

Explanation

An I/O error occurred when connecting to the specified host and port.

Action

Verify network connectivity to the specified host and port.

GEN0335E Unable to send a test email message to *server name*.

Explanation

IBM Spectrum Control cannot send a test email to the specified SMTP server.

Action

Verify network connectivity to the specified mail server. Ensure the server is active and the port is the unsecured mail server port.

GEN0400I Probe aborted.

Explanation

The probe was aborted.

GEN0400E Probe completed with errors.

Explanation

The probe completed with errors. The job and server logs contain error messages detailing the errors.

Action

Check the job and server logs for additional information about the errors. If unable to determine the cause of the errors, contact IBM customer technical support.

Related reference

- [Getting support](#)
-

GEN0401I Probe completed successfully.

Explanation

The probe completed successfully.

Action

No action is required. If you want, you can examine the job log for further details.

GEN0402W Probe completed with warnings.

Explanation

The probe completed with warnings.

Action

Check the job and server logs for details about the warnings.

GEN0403E Unable to retrieve probe definition.

Explanation

The server was unable to retrieve the probe definition from the repository.

Action

Examine the job and server logs and determine what caused the error. If unable to determine the cause of the error, contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0404I Probe started.

Explanation

The probe has been started.

GEN0405E No computers or storage subsystems to probe.

Explanation

The probe definition did not specify any valid element to be probed.

Action

Specify at least one valid element to be probed.

GEN0406E A probe of the computer or hypervisor *computer name* is already in progress.

Explanation

Another probe of the same computer or hypervisor is already in progress, so the new probe cannot be started.

Action

Ensure that the previous probe has completed before starting a new probe for the same computer or hypervisor.

GEN0407E Probe of storage subsystem *subsystem name (subsystem alias)* already in progress.

Explanation

Another probe of the same storage subsystem is already in progress, so the new probe cannot be started.

Action

Ensure that the previous probe has completed before starting a new probe for the same storage subsystem.

GEN0408W Probe of fabric *fabric WWN* already in progress.

Explanation

Another probe of the same fabric is already in progress, so the new probe cannot be started.

Action

Ensure that the previous probe has completed before starting a new probe for the same fabric.

GEN0409E Batch Report job *report name* already in progress.

Explanation

Another job for the same batch report is already in progress, so the new one cannot be started.

Action

Ensure that the previous job has completed before starting a new one for the same batch report.

GEN0410E Probe of switch *switch WWN* already in progress.

Explanation

Another probe of the same switch is already in progress, so the new probe cannot be started.

Action

Ensure that the previous probe has completed before starting a new probe for the same switch.

GEN0557E An OS error occurred.

Explanation

An unexpected OS error occurred. Error information will be logged.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN0558E An IO error occurred.

Explanation

An unexpected IO error occurred. Error information will be logged.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN1030E *product name* is not licensed on computer *computer name*.

Explanation

The specified product is not licensed on the specified computer.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN1034I GeneralException message text follows:*error message*.

Explanation

An informational message displaying the General Exception message.

GEN1035W *storage subsystem name* is no longer monitored.

Explanation

The specified storage subsystem is no longer monitored.

Action

Verify that the SMI-S provider where the storage subsystem is defined has correctly been saved in the IBM Spectrum Control SMI-S provider Login Administration panel. If necessary, verify that the login information is correct.

GEN6013E OS Error *errno*: *error description*

Explanation

The installer or agent has encountered an OS specific error.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN6014W OS Error *errno*

Explanation

The installer or agent has encountered an OS specific error.

Action

Contact IBM customer technical support.

GEN6015E Unable to load JAVA.DLL

Explanation

The installer or agent failed to load the JAVA.DLL library.

Action

Contact IBM customer technical support.

Related reference

- [🔗 Getting support](#)

GEN6016E Unable to find symbol *module symbol* in JAVA.DLL

Explanation

The installer or agent failed to find the symbol in the JAVA.DLL library.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN6017E Executable file is null

Explanation

The path to the executable file is null.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN6018E *stat(file path)* failed.

Explanation

The file status `stat()` command failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN6019E *chmod(file path)* failed.

Explanation

The change of permissions mode `chmod()` command failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

GEN5001W Not all statistics could be saved due to unlicensed computer(s). See TPCD log for details.

Explanation

Not all statistics could be saved due to unlicensed computers.

GEN7111E Unable to retrieve data from the repository database.

Explanation

IBM Spectrum Control cannot retrieve data from the repository database. The repository database contains information that was collected about the storage resource.

Action

Ensure that the repository database, Data server, and Device server are running. Verify that the local area network is available. Check the product log files for error messages that might help determine the problem. See the product information center to view the locations of these log files. Try the action again. If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

GEN7112E Failed to retrieve e-mail report from (*report URL*).

Explanation

The report could not be downloaded from the specified URL.

Action

Check that the URL is accessible. If necessary, start the web server and retry.

GEN7113E Failed to retrieve e-mail report from the web server.

Explanation

The report could not be downloaded from the web server due to an unknown error.

Action

Check that the web server is up. If necessary, start the web server and retry.

If this does not resolve the issue then contact IBM support.

Related reference

- [Getting support](#)

GEN7114E The user and password combination for the outgoing e-mail server is invalid.

Explanation

Failed to authenticate with the outgoing e-mail server..

Action

Update the user and password in the e-mail server configuration panel.

If this does not resolve the issue then contact IBM support.

Related reference

- [Getting support](#)

GEN7115E Can't send the attached "*reportName*" report with the ID *reportId*. The size of the attachment (*totalSize* KB) exceeds the maximum size for attachments that was set on your email server. [Learn More link](#)

Explanation

The size of the message exceeds the physical limit or the administrative limit that is set for messages on the email server.

Action

Change the size limit on the email server that you configured to send reports.

GEN7116E The email address that was configured for replies to alert notifications and reports, *reply_to_address_value*, was not accepted by the email server, *server*.

Explanation

The email address that was configured for replies to alert notifications and reports was rejected.

Action

Check the documentation for the email server that you use to send emails. For example, your email server might expect that the user name that you configure in IBM Spectrum Control to authenticate with your email server is an email address and that you must use the same address for the user name and the reply-to address.

GPC - Performance User Interface messages

- [GPC000001E One or more of the selected profiles is a system-defined profile and cannot be deleted. Deselect profile and try again.](#)
- [GPC000002E The selected profile profile is a system-defined profile and cannot be updated.](#)
- [GPC000003E There is no performance data collected for any of the storage subsystems, create workload profile Wizard can not continue.](#)
- [GPC000150W Analyzing all the volume performance data may take a long time to complete.](#)
- [GPC000200E Based on the current choices, no devices can be included in the analysis.](#)
- [GPC000201E There is no ESS performance data collected, Volume Performance Advisor Wizard can not continue.](#)
- [GPC000350I Volume size may round up to multiple of 100 MB.](#)
- [GPC00400E Subsystem type for type identifier does not match the subsystem type in selected list.](#)
- [GPC00401E Duration value is not specified.](#)
- [GPC00800E The value entered for the number of rows is incorrect.](#)
- [GPC00801E The selected metrics do not have same unit type. Charts only display metrics with the same unit type.](#)
- [GPC00802E The selected charting metrics do not have same chart type. Chart and History Chart options cannot be mixed.](#)
- [GPC00803E Performance reports can only be saved with 300 or less explicitly specified components. To save the report either select all components or select no more than 300.](#)
- [GPC00804E Please enter a relative time higher than zero.](#)
- [GPC00950I There is no data to be charted.](#)
- [GPC00951I No metrics were selected.](#)
- [GPC00952I No drill-up available for this component.](#)
- [GPC00953I None of the selected metrics apply for all the report components.](#)
- [GPC00954I None of the selected metrics apply for all selected components.](#)
- [GPC00955I None of the selected metrics apply for all the report components or for all selected components.](#)
- [GPC00956I Subsystem port constraint violations do not support the affected volumes report.](#)
- [GPC00957I The maximum limit of max no of rows displayable rows for a performance report has been reached. If a larger report is required rerun the report as a batch report, alternatively redefine the report to return less data. The amount of report data can be reduced by using a smaller time window, working with less components, applying filters when defining the report or using the aggregated hourly or daily report data rather than the sample data.](#)
- [GPC00958I No drill-down available for this component.](#)
- [GPC50000E Fabric type for fabric type identifier does not match the fabric type in selected list.](#)

GPC000001E One or more of the selected profiles is a system-defined profile and cannot be deleted. Deselect *profile* and try again.

Explanation

Action

GPC000002E The selected profile *profile* is a system-defined profile and cannot be updated.

Explanation

Action

GPC000003E There is no performance data collected for any of the storage subsystems, create workload profile Wizard can not continue.

Explanation

Action

GPC000150W Analyzing all the volume performance data may take a long time to complete.

Explanation

Action

GPC000200E Based on the current choices, no devices can be included in the analysis.

Explanation

After excluding the subsystems with missing host ports, there are no devices available for analysis. A new selection of subsystems and host port combinations is required.

Action

GPC000201E There is no ESS performance data collected, Volume Performance Advisor Wizard can not continue.

Explanation

Action

GPC000350I Volume size may round up to multiple of 100 MB.

Explanation

Action

GPC00400E Subsystem type for *type identifier* does not match the subsystem type in selected list.

Explanation

The device selected in the source list does not have the same subsystem type in the selected list.

Action

GPC00401E Duration value is not specified.

Explanation

Duration value is not specified. Please specify a value.

Action

GPC00800E The value entered for the number of rows is incorrect.

Explanation

The user entered a non-integer value for the number of rows.

Action

GPC00801E The selected metrics do not have same unit type. Charts only display metrics with the same unit type.

Explanation

The charts display a single Y-Axis listing the metric unit types, hence the charts only display the metrics of the same unit type.

Action

GPC00802E The selected charting metrics do not have same chart type. Chart and History Chart options cannot be mixed.

Explanation

There is a choice to display batch chart or batch history chart reports, options for both chart types cannot be mixed.

Action

GPC00803E Performance reports can only be saved with 300 or less explicitly specified components. To save the report either select all components or select no more than 300.

Explanation

The database is limited in the number of components a saved report can contain. The limit is either 300 user-specified components or all components. All components can be saved in the report since a flag is used to indicate the condition rather than an entry for each component.

Action

GPC00804E Please enter a relative time higher than zero.

Explanation

The relative time in days to display the historic performance data should be higher than zero.

Action

GPC00950I There is no data to be charted.

Explanation

There is no data available for the requested chart.

Action

GPC00951I No metrics were selected.

Explanation

No metrics were selected when this report was generated.

Action

GPC00952I No drill-up available for this component.

Explanation

Drill-up is not available for this component.

Action

GPC00953I None of the selected metrics apply for all the report components.

Explanation

A metric can be charted only if it applies to each component in the report.

Action

GPC00954I None of the selected metrics apply for all selected components.

Explanation

A metric can be charted only if it applies to each selected component.

Action

GPC00955I None of the selected metrics apply for all the report components or for all selected components.

Explanation

A metric can be charted only if it applies for each component in the report or for each selected component.

Action

GPC00956I Subsystem port constraint violations do not support the affected volumes report.

Explanation

Affected volumes reports are not generated for metrics associated with the subsystem port component.

Action

GPC00957I The maximum limit of *max no of rows* displayable rows for a performance report has been reached. If a larger report is required rerun the report as a batch report, alternatively redefine the report to return less data. The amount of report data can be reduced by using a smaller time window, working with less components, applying filters when defining the report or using the aggregated hourly or daily report data rather than the sample data.

Explanation

The IBM Spectrum Control GUI has to limit the number of rows in a performance report to avoid heapdumps caused by large reports.

Action

If the user wants to view more data than the default 100000 rows in the GUI, the user can modify data server's TPCD.config parameter, pmReportRowLimit. To avoid the GUI running out of memory the GUI's max heapsize should also be increased.

GPC00958I No drill-down available for this component.

Explanation

Drill-down is not available for this component.

Action

GPC50000E Fabric type for *fabric type identifier* does not match the fabric type in selected list.

Explanation

The device selected in the source list does not have the same subsystem type in the selected list.

Action

HWNAS - Agentless Server messages

- [HWNAS0001I Successfully created server server.](#)
- [HWNAS0002I Successfully deleted server server.](#)
- [HWNAS0003E The host name or IP address {0} is not valid.](#)
- [HWNAS0004E Cannot add port portWWPN because it belongs to server serverName.](#)
- [HWNAS0005E Cannot add port portWWPN because it belongs to switch switchName.](#)
- [HWNAS0006E Cannot add port portWWPN because it belongs to storage system storageSystemName.](#)
- [HWNAS0007W Server serverName was not created because it exists already.](#)
- [HWNAS0008I Successfully created mergeServerName server by merging numberOfServers servers.](#)
- [HWNAS0009I Successfully separated serverName server into numberOfServers individual servers.](#)
- [HWNAS0010E The serverId agentless server that you selected does not exist.](#)
- [HWNAS0011I You cannot separate the serverName agentless server because it is not based on storage system host connections.](#)
- [HWNAS0012I You cannot separate the serverName agentless server because it is already defined on the smallest possible separation boundary.](#)
- [HWNAS0013I You cannot merge the selected agentless servers into the serverName agentless server because they are not all based on storage system host connections.](#)

HWNAS0001I Successfully created server server.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNAS0002I Successfully deleted server server.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNAS0003E The host name or IP address {0} is not valid.

Explanation

The IP address or host name that was entered for the device is not valid.

Action

Make sure that the IP address and host name are valid for the device that is being added. Reenter the IP address or host name and click Add again.

HWNAS0004E Cannot add port portWWPN because it belongs to server serverName.

Explanation

The specified port WWPN belongs to another server and cannot be added to the server.

Action

Enter a port WWPN that does not belong to another server.

HWNAS0005E Cannot add port *portWWPN* because it belongs to switch *switchName*.

Explanation

The specified port WWPN is a switch port and cannot be added to the server.

Action

Enter a port WWPN that does not belong to a switch.

HWNAS0006E Cannot add port *portWWPN* because it belongs to storage system *storageSystemName*.

Explanation

The specified port WWPN is a storage system port and cannot be added to the server.

Action

Enter a port WWPN that does not belong to a storage system.

HWNAS0007W Server *serverName* was not created because it exists already.

Explanation

The server was not created because a server with the same host name or IP address exists already.

Action

Enter a host name or IP address for a server that is not already monitored.

HWNAS0008I Successfully created *mergeServerName* server by merging *numberOfServers* servers.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNAS0009I Successfully separated *serverName* server into *numberOfServers* individual servers.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNAS0010E The *serverId* agentless server that you selected does not exist.

Explanation

Either the server ID is invalid or it refers to a resource that is not an agentless server. For example the server might be a server with a Storage Resource agent deployed.

Action

Select a valid agentless server.

HWNAS0011I You cannot separate the *serverName* agentless server because it is not based on storage system host connections.

Explanation

The agentless server has no assigned WWPNs, or the assigned WWPNs are not associated with any storage system host connections. For example, the agentless server might represent a virtual machine that was discovered during a hypervisor probe. You cannot separate these agentless servers into multiple individual servers.

Action

Select an agentless server that represents two or more storage system host connections.

HWNAS0012I You cannot separate the *serverName* agentless server because it is already defined on the smallest possible separation boundary.

Explanation

The separation of an agentless server occurs on boundaries that are identified by storage system host connections that have sets of WWPNs that do not overlap. The agentless server that you selected is either based on a single host connection and therefore cannot be further subdivided, or the agentless server is based on multiple host connections that all have overlapping sets of WWPNs assigned.

Action

If the agentless server is based on multiple host connections and you want to separate it, use the GUI or CLI for the storage system to modify the host connection definitions. Remove any unintended overlap between sets of WWPNs for the storage system host connections. After you remove the overlap, rerun the probe on those storage systems. After the probes complete, try to separate the agentless server again.

You cannot eliminate any overlap that is caused by HBAs. HBAs are discovered during the fabric probe, and identify which WWPNs cannot be separated. That is, the HBAs identify which WWPNs are known to belong to the same interface card in a server computer.

HWNAS0013I You cannot merge the selected agentless servers into the *serverName* agentless server because they are not all based on storage system host connections.

Explanation

One or more of the agentless servers is not associated with any storage system host connections. For example, the agentless server might represent a virtual machine that was discovered during a hypervisor probe. You cannot merge these agentless servers into a single server.

Action

Select agentless servers that are based on storage system host connections.

HWNAL - Single sign-on service messages

- [HWNAL0001E A connection with the IBM Spectrum Control Device Server, \(Device Server IP, could not be established for authentication.](#)

- [HWNAU0002E A connection with the LDAP or Active Directory server, \(LDAP or Active Directory Server IP \) , could not be made for authentication.](#)
- [HWNAU0003E Authentication of the Single Sign-On token failed. Provide your username and password to attempt a re-authentication.](#)
- [HWNAU0004E The Single Sign-On token has expired. To re-authenticate the token, please enter your user name / password.](#)
- [HWNAU0005E Creation of the Single Sign-On token failed due to an username that is not valid. Enter your username and password and try again.](#)
- [HWNAU0006E Creation of the Single Sign-On token failed due to a password that is not valid. Enter your username and password and try again.](#)
- [HWNAU0007E Authentication failed due to an username or password that is not valid. Enter your username and password and try again.](#)
- [HWNAU0008I Single Sign On Service started successfully.](#)
- [HWNAU0009I The Single Sign On Service has shutdown.](#)
- [HWNAU0010E An error occurred retrieving the Single Sign-On token from the private credentials.](#)
- [HWNAU0011E An error occurred retrieving the Single Sign-On token from the public credentials.](#)
- [HWNAU0012E An error occurred when attempting to decode the authentication token.](#)
- [HWNAU0013E An error occurred when attempting to encode the authentication token.](#)
- [HWNAU0014E An error occurred while translating the user's credentials into a Single Sign-On token.](#)
- [HWNAU0015E An unknown error occurred while authenticating to the WebSphere login module.](#)
- [HWNAU0016E An error occurred while registering SsoConfigChangeListener with TIP.](#)
- [HWNAU0017E An error occurred while unregistering SsoConfigChangeListener from TIP.](#)
- [HWNAU0018E The Web server appears to be down and cannot be used for authentication. It is still possible to perform OS user authentication against the device server, however since the Web server is down the IBM Spectrum Control functionality will be limited. Among the limitations is the inability to perform SSO to other applications that rely on the presence of a lightweight third party authentication token. To proceed enter a local OS username with administrative privileges and password.](#)
- [HWNAU0019E An unknown error occurred while authenticating with Web server.](#)

HWNAU0001E A connection with the IBM Spectrum Control Device Server, (Device Server IP, could not be established for authentication.

Explanation

A connection to the IBM Spectrum Control Device server could not be established for the authentication request.

Action

Make sure the IBM Spectrum Control Device Server is active and has IP connectivity with the IBM Spectrum Control server.

HWNAU0002E A connection with the LDAP or Active Directory server, (LDAP or Active Directory Server IP) , could not be made for authentication.

Explanation

A connection to the LDAP or Active Directory server could not be made for the authentication request.

Action

Make sure the LDAP or Active Directory server is active and has IP connectivity with the IBM Spectrum Control Device Server.

HWNAU0003E Authentication of the Single Sign-On token failed. Provide your username and password to attempt a re-authentication.

Explanation

Authentication of the Single Sign-On token failed.

Action

Make sure the Single Sign-On token was created within the same domain as the IBM Spectrum Control server.

HWNAU0004E The Single Sign-On token has expired. To re-authenticate the token, please enter your user name / password.

Explanation

The Single Sign-On token expired. Enter your password to resume the Single Sign-On login session.

Action

Provide the Single Sign-On token's password in order for it to reauthenticate.

HWNAU0005E Creation of the Single Sign-On token failed due to an username that is not valid. Enter your username and password and try again.

Explanation

Creation of the Single Sign-On token failed due to an username that is not valid.

Action

Either make sure the username and password are valid or the user is configured to the LDAP or Active Directory server.

HWNAU0006E Creation of the Single Sign-On token failed due to a password that is not valid. Enter your username and password and try again.

Explanation

Creation of the Single Sign-On token failed due to a password that is not valid.

Action

Make sure the password of the user is valid.

HWNAU0007E Authentication failed due to an username or password that is not valid. Enter your username and password and try again.

Explanation

Authentication failed due to an username or password that is not valid.

Action

Make sure the username and password are valid.

HWNAU0008I Single Sign On Service started successfully.

Explanation

The Single Sign On Service has started successfully.

HWNAU0009I The Single Sign On Service has shutdown.

Explanation

The Single Sign On Service has been shutdown.

HWNAU0010E An error occurred retrieving the Single Sign-On token from the private credentials.

Explanation

An error occurred retrieving the Single Sign-On token from the private credentials.

HWNAU0011E An error occurred retrieving the Single Sign-On token from the public credentials.

Explanation

An error occurred retrieving the Single Sign-On token from the public credentials.

HWNAU0012E An error occurred when attempting to decode the authentication token.

Explanation

An error occurred when attempting to decode the authentication token.

HWNAU0013E An error occurred when attempting to encode the authentication token.

Explanation

An error occurred when attempting to encode the authentication token.

HWNAU0014E An error occurred while translating the user's credentials into a Single Sign-On token.

Explanation

An error occurred while translating the user's credentials into a Single Sign-On token.

HWNAU0015E An unknown error occurred while authenticating to the WebSphere login module.

Explanation

An unknown error occurred while authenticating to the WebSphere login module.

HWNAU0016E An error occurred while registering SsoConfigChangeListener with TIP.

Explanation

An error occurred while registering SsoConfigChangeListener with TIP.

HWNAU0017E An error occurred while unregistering SsoConfigChangeListener from TIP.

Explanation

An error occurred while unregistering SsoConfigChangeListener from TIP.

HWNAU0018E The Web server appears to be down and cannot be used for authentication. It is still possible to perform OS user authentication against the device server, however since the Web server is down the IBM Spectrum Control functionality will be limited. Among the limitations is the inability to perform SSO to other applications that rely on the presence of a lightweight third party authentication token. To proceed enter a local OS username with administrative privileges and password.

Explanation

Authentication and Single Sign On (SSO) functionality is primarily obtained through the Web server. Without the Web server, IBM Spectrum Control can only authenticate the user with no SSO capability.

Action

Make sure the Web Server is active and has IP connectivity with the IBM Spectrum Control server.

HWNAU0019E An unknown error occurred while authenticating with Web server.

Explanation

An unknown error occurred while authenticating with Web Server.

HWNCA - Storage multiple access control messages

- [HWNCA0001E The required parameter `name` parameter is missing.](#)
- [HWNCA0002E The delivery unit type `delivery_unit_type` is not supported.](#)
- [HWNCA0003E The delivery unit type, `delivery_unit_type`, does not match the capacity pool type, `capacity_pool_type`.](#)
- [HWNCA0004E The service class `service_class_name` is not supported by the capacity pool.](#)
- [HWNCA0005E IBM Spectrum Control cannot provide the requested delivery unit size GiB of capacity that also satisfies the requirements of the service class.](#)
- [HWNCA0006E A delivery unit `delivery_unit_name` and service instance ID `service_instance_id` already exists.](#)
- [HWNCA0007I Normal isolated file storage.](#)
- [HWNCA0008I Enhanced isolated file storage.](#)
- [HWNCA0009E The capacity pool with ID `capacity_pool_id` is still in use and cannot be deleted.](#)
- [HWNCA0010E The service class `service_class_name` has no recipe for deleting delivery unit `delivery_unit_name`.](#)
- [HWNCA0011E The service class `service_class_name` cannot modify the delivery unit `delivery_unit_name`.](#)
- [HWNCA0012E The capacity of the delivery unit with ID `delivery_unit_id` cannot be modified after delivery unit is created.](#)
- [HWNCA0013E Delivery unit with ID `delivery_unit_id` cannot be modified or deleted because it is in the processing state.](#)
- [HWNCA0014E Delivery unit with ID `delivery_unit_id` cannot be modified because it has a completion state error.](#)
- [HWNCA0015E The principal `principal_name` was used in multiple CIFS access control list \(ACL\) entries.](#)
- [HWNCA0016I A standard object can handle object-type delivery units, such as CDMI containers.](#)
- [HWNCA0017I You can use the StandardBlock service class to create, modify, and delete delivery units that are block based, such as volume containers.](#)
- [HWNCA0018I High Availability, Sequential.](#)
- [HWNCA0019I High Availability, Short response time, Transactional.](#)
- [HWNCA0020E No provisioning profile was found for recipe name.](#)
- [HWNCA0021E The WWPN `WWPN` does not belong to a server or hypervisor managed by IBM Spectrum Control.](#)
- [HWNCA0022E Device Selection failed with the following error: Planner Message](#)
- [HWNCA0023E Creation of DeliveryUnit failed with the following error: Planner Message](#)
- [HWNCA0025I Remote Caching allows the data written to one delivery unit to be pushed to another remote located delivery unit.](#)
- [HWNCA0026E Deletion of DeliveryUnit failed with the following error: Planner Message](#)
- [HWNCA0027I Highest-performing storage for mission-critical applications.](#)

- [HWNCA0028I High-performing storage for applications in production.](#)
- [HWNCA0029I Standard storage for non-mission-critical applications.](#)
- [HWNCA0030E No host port WWPNs were found for server host name.](#)
- [HWNCA0031I The standard root object class enforces the use of dedicated filesets for associated object-type delivery units, such as CDMI containers.](#)
- [HWNCA0032E A Storage Resource agent is not up and running for the host name server.](#)
- [HWNCA0033E A service class with name service_class_name does not exist for the specified delivery unit type.](#)
- [HWNCA0034E IBM Spectrum Control is unable to create or modify the service class because one or more of the capacity pools to which it was restricted were removed.](#)
- [HWNCA0035E Delivery unit with ID delivery_unit_id cannot be modified or deleted because it has a completion status of planned or change_planned.](#)
- [HWNCA0036E More than numResults results found. Modify your search pattern.](#)
- [HWNCA0037I Start the creation of an object container with name containerName.](#)
- [HWNCA0038I The creation of the object container with name containerName is finished.](#)
- [HWNCA0039I Start the deletion of an object container with name containerName.](#)
- [HWNCA0040I The deletion of the object container with name containerName is finished.](#)
- [HWNCA0041I Start the modification of an object container with name containerName.](#)
- [HWNCA0042I The modification of the object container with name containerName is finished.](#)
- [HWNCA0043I This task was already deleted.](#)
- [HWNCA0044E The requested operation requires SMAC API client version version or later.](#)
- [HWNCA0045E The plan with scheduleID scheduleID can not be deleted because the delivery unit creation is already processing state.](#)
- [HWNCA0046E The specified value deletionType is not a valid plan deletion type.](#)
- [HWNCA0047E Host with host name host is not known by IBM Spectrum Control.](#)
- [HWNCA0048E Filer Storage ID filerId for given fileSystem is not found.](#)
- [HWNCA0049W IBM Spectrum Control did not change any Fibre Channel zones. After the provisioning operation completes, verify the Fibre Channel connectivity between any involved host and the storage subsystem that contains the storage volume. If necessary, change the zoning configuration.](#)
- [HWNCA0050W IBM Spectrum Control did not enable the multipath policy for the host or hosts to which the volume was reassigned. You must enable the multipath policy after the provisioning operation completes.](#)
- [HWNCA0051E Incorrect volume name volumeName. The name can contain letters, numbers, dashes \(-\) and underscores \(_\). However, the name cannot start with a number, dash, or the reserved words vdisk or volume.](#)
- [HWNCA0052E One or more of the requested volume names are too long for the following storage systems: StorageSubsystemsList. The maximum volume name length that is allowed is MaxAllowedNameLength characters, and a volume name of LargestNameLength characters was requested.](#)
- [HWNCA0053E The service class scName does not support provisioning from any available storage.](#)
- [HWNCA0054E The volume volumeName already exists on the storage system storageSystemName. Request a different volume name.](#)
- [HWNCA0055E The share shareName already exists on the storage system storageSystemName. Request a different share name.](#)
- [HWNCA0056W Zoning cannot be configured because the connected fabrics are not monitored.](#)
- [HWNCA0057W The multipath policy cannot be set on the following hypervisors: HostList.](#)
- [HWNCA0058W The multipath policy cannot be set on the following agentless servers: HostList.](#)
- [HWNCA0059W The multipath policy cannot be set on the following hosts: HostList.](#)
- [HWNCA0060E Storage volume can not be assigned to servers that are running different operating systems.](#)
- [HWNCA0061E The required resources cannot be reserved because too many provisioning tasks are currently being created. Try again in a few minutes.](#)
- [HWNCA0062E The operation did not complete because the specified service class or capacity pool no longer exists.](#)
- [HWNCA0063E The service class ScName that is specified by this provisioning task does not exist.](#)
- [HWNCA0064E The current SMAC API client is version current_version. The requested operation requires SMAC API client version required_version or later.](#)
- [HWNCA0065E The delivery unit specified has snapshots and could not be deleted.](#)
- [HWNCA0066E The WWPN WWPN already belongs to a server or hypervisor managed by IBM Spectrum Control. Was expected an unknown WWPN based on previous Hosts or WWPNs passed in.](#)
- [HWNCA0067E The host host already belongs to a server or hypervisor managed by IBM Spectrum Control. Was expected an unknown Host based on previous Hosts or WWPNs passed in.](#)
- [HWNCA0068E The list of hosts was expected to contain just one or no element because the WWPNs passed in were all unknown. The size of the host list is size.](#)
- [HWNCA0069E No unknown WWPNs were specified for this provisioning operation but an unknown Host was specified.](#)
- [HWNCA0070E The expansion of the capacity is not supported for the delivery unit with ID delivery_unit_id.](#)
- [HWNCA0071E The reduction of the capacity is not supported for the delivery unit with ID delivery_unit_id.](#)

HWNCA0001E The required *parameter_name* parameter is missing.

Explanation

A required parameter was not specified.

Action

Try the command again, specifying all parameters.

HWNCA0002E The delivery unit type *delivery_unit_type* is not supported.

Explanation

The specified delivery unit type is not supported by the Storage Management API for Cloud.

Action

HWNCA0003E The delivery unit type, *delivery_unit_type*, does not match the capacity pool type, *capacity_pool_type*.

Explanation

The specified delivery unit type does not match the type of the capacity pool that is in use.

Action

Specify a different delivery unit type or use a different capacity pool to allocate the delivery unit.

HWNCA0004E The service class *service_class_name* is not supported by the capacity pool.

Explanation

The capacity pool contains no resources that support the requested service class.

Action

Add resources to this capacity pool that support the service class, or specify a different capacity pool for this service class.

HWNCA0005E IBM Spectrum Control cannot provide the requested *delivery_unit_size* GiB of capacity that also satisfies the requirements of the service class.

Explanation

There is not enough storage available that also satisfies the requirements of the specified service class. If the provisioning request is constrained to a capacity pool, the capacity pool does not contain enough storage that also satisfies the requirements of the specified service class.

Action

Review the service class properties to determine if the service class requirements are too restrictive, or if additional storage should be added. Depending on your findings, complete one of the following steps:

- Specify a different service class or modify the properties of the service class.
- Change the provisioning constraints. If you are limiting provisioning to a capacity pool, specify a different capacity pool or allow provisioning from any available storage.
- Add storage that satisfies the requirements of the service class.
- Check the msgPlanner.log from <IBM Spectrum Control folder>/device/log for more information about the provisioning task.

HWNCA0006E A delivery unit *delivery_unit_name* and service instance ID *service_instance_ID* already exists.

Explanation

The delivery unit already exists.

Action

Choose a different name for the delivery unit or assign the delivery unit to a different service instance.

HWNCA0007I Normal isolated file storage.

Explanation

Normal isolation allows the shared use of storage resources for NAS exports.

Action

HWNCA0008I Enhanced isolated file storage.

Explanation

Enhanced isolation enforces the use of dedicated storage resources for NAS exports.

Action

HWNCA0009E The capacity pool with ID *capacity_pool_id* is still in use and cannot be deleted.

Explanation

At least one delivery unit allocates storage from the capacity pool. The capacity pool cannot be deleted until all delivery units that allocate storage from it have been deleted.

Action

Delete all delivery units that allocate storage from this capacity pool. Try to delete the capacity pool again.

HWNCA0010E The service class *service_class_name* has no recipe for deleting delivery unit *delivery_unit_name*.

Explanation

The delivery unit could not be deleted because the recipe for this device type was not found.

Action

Add a recipe for this service class that supports the device type that the delivery unit was created on.

HWNCA0011E The service class *service_class_name* cannot modify the delivery unit *delivery_unit_name*.

Explanation

The delivery unit could not be modified because no recipe for this device type was found.

Action

Add a recipe for this service class that supports the device type that the delivery unit was created on.

HWNCA0012E The capacity of the delivery unit with ID *delivery_unit_id* cannot be modified after delivery unit is created.

Explanation

IBM Spectrum Control was unable to modify the delivery unit's capacity, because the corresponding recipe does not support this operation.

Action

Replace the existing delivery unit with a new one that has the needed capacity. Alternatively, create an additional delivery unit to extend the available capacity of this delivery unit.

HWNCA0013E Delivery unit with ID *delivery_unit_id* cannot be modified or deleted because it is in the processing state.

Explanation

The creation process of this delivery unit is not complete. You cannot delete or modify a delivery unit that is in a processing state.

Action

Wait until the creation process is complete. Check the attributes of the delivery unit to see the status of the creation process.

HWNCA0014E Delivery unit with ID *delivery_unit_id* cannot be modified because it has a completion state error.

Explanation

An error occurred during the creation of this delivery unit. You cannot modify a delivery unit that has creation errors.

Action

Delete and then create the delivery unit again.

HWNCA0015E The principal *principal_name* was used in multiple CIFS access control list (ACL) entries.

Explanation

CIFS ACLs do not support the specification of the same principal in multiple entries.

Action

Remove duplicate ACL entries.

HWNCA0016I A standard object can handle object-type delivery units, such as CDMI containers.

Explanation

This service class allows the creation, modification, and deletion of object-type delivery units, such as CDMI containers.

Action

No further action is required.

HWNCA0017I You can use the StandardBlock service class to create, modify, and delete delivery units that are block based, such as volume containers.

Explanation

This message is for informational purposes only.

Action

No further action is required.

HWNCA0018I High Availability, Sequential.

Explanation

This service class allows the creation, modification, and deletion of block-type delivery units, such as volume containers using a default RAID5 provisioning profile.

Action

No further action is required.

HWNCA0019I High Availability, Short response time, Transactional.

Explanation

This service class allows the creation, modification, and deletion of block-type delivery units, such as volume containers with a default RAID10 provisioning profile.

Action

No further action is required.

HWNCA0020E No provisioning profile was found for *recipe name*.

Explanation

No provisioning profile was found for the PlannerEnabledBlockRecipe.

Action

Specify a provisioning profile and try again.

HWNCA0021E The WWPN *WWPN* does not belong to a server or hypervisor managed by IBM Spectrum Control.

Explanation

The Host ports/WWPNs must belong to a server or hypervisor managed by IBM Spectrum Control.

Action

Probe the server or select another WWPN and try again.

HWNCA0022E Device Selection failed with the following error: *Planner Message*

Explanation

The Device Selection failed.

Action

Perform the actions specified by the Planner and try again.

HWNCA0023E Creation of DeliveryUnit failed with the following error: *Planner Message*

Explanation

Creation of DeliveryUnit failed.

Action

Perform the actions specified by the Planner and try again.

HWNCA0025I Remote Caching allows the data written to one delivery unit to be pushed to another remote located delivery unit.

Explanation

Delivery units created with this service class can be configured so that data written to them is automatically pushed to another, potentially offsite, storage location.

Action

HWNCA0026E Deletion of DeliveryUnit failed with the following error: *Planner Message*

Explanation

Deletion of DeliveryUnit failed.

Action

HWNCA0027I Highest-performing storage for mission-critical applications.

Explanation

Storage volumes that are associated with this service class are placed on storage resources that provide the highest performance characteristics.

Action

HWNCA0028I High-performing storage for applications in production.

Explanation

Storage volumes that are associated with this service class are placed on storage resources that provide high performance characteristics.

Action

HWNCA0029I Standard storage for non-mission-critical applications.

Explanation

Storage volumes that are associated with this service class are placed on storage resources that provide standard performance characteristics.

Action

HWNCA0030E No host port WWPNs were found for server *host name*.

Explanation

To contain World Wide Port Names (WWPNs), a server must be managed by a Storage Resource agent or an ESX server.

Action

Ensure that the server is managed by a Storage Resource agent or an ESX server. After the server is managed, run a probe to collect data about its WWPNs. Try to create or modify the Delivery Unit again.

HWNCA0031I The standard root object class enforces the use of dedicated filesets for associated object-type delivery units, such as CDMI containers.

Explanation

You can use this service class to create, modify and delete isolated object-type delivery units, such as CDMI containers. On creation, a dedicated fileset uses available free storage resources as needed.

Action

No further action is required.

HWNCA0032E A Storage Resource agent is not up and running for the *host name server*.

Explanation

To assign World Wide Port Names (WWPNs), the Storage Resource agent that is managing a server must be up and running.

Action

Ensure that the server is managed by a Storage Resource agent and start the agent. Try to create or modify the Delivery Unit again.

HWNCA0033E A service class with name *service_class_name* does not exist for the specified delivery unit type.

Explanation

The service class could not be found.

Action

Ensure that such a service class exists and check that its name was spelled correctly.

HWNCA0034E IBM Spectrum Control is unable to create or modify the service class because one or more of the capacity pools to which it was restricted were removed.

Explanation

The service class that you attempted to create or modify defines storage constraints that restrict future provisioning to a specific set of capacity pools. While the service class creation or modification was in progress, one or more of the specified capacity pools were removed. Because one or more of the specified capacity pools no longer exist, IBM Spectrum Control cannot create or modify the service class.

Action

In the Create Service Class wizard or service class properties notebook, go to the list of capacity pools. The capacity pools that were deleted are no longer listed. Define the storage constraints from the new list of capacity pools.

HWNCA0035E Delivery unit with ID *delivery_unit_id* cannot be modified or deleted because it has a completion status of planned or change_planned.

Explanation

You cannot modify or delete a delivery unit that in planned or change planned state.

Action

Execute or delete the existing plan to allow delivery unit modification or deletion.

HWNCA0036E More than *numResults* results found. Modify your search pattern.

Explanation

The number of users that were found with the provided search pattern exceeds the limit.

Action

Modify your search pattern and try again.

HWNCA0037I Start the creation of an object container with name *containerName*.

Explanation

The operation to start the creation of object container is started.

Action

No further action is required.

HWNCA0038I The creation of the object container with name *containerName* is finished.

Explanation

The operation to start the creation of object container is completed.

Action

No further action is required.

HWNCA0039I Start the deletion of an object container with name *containerName*.

Explanation

The operation to delete the object container is started.

Action

No further action is required.

HWNCA0040I The deletion of the object container with name *containerName* is finished.

Explanation

The operation to delete the object container is completed.

Action

No further action is required.

HWNCA0041I Start the modification of an object container with name *containerName*.

Explanation

The operation to modify the object container is started.

Action

No further action is required.

HWNCA0042I The modification of the object container with name *containerName* is finished.

Explanation

The operation to modify the object container is completed.

Action

No further action is required.

HWNCA0043I This task was already deleted.

Explanation

You are attempting to execute, schedule, or delete a task that was already deleted. The task was possibly deleted by another user.

Action

No action is required.

HWNCA0044E The requested operation requires SMAC API client version *version* or later.

Explanation

Some SMAC API operations require the most recent version of the SMAC client. This prevents unexpected behavior by older SMAC API clients.

Action

Update the SMAC API client to the version specified in this message, or to a more recent version. For best results, the SMAC API client should be at the same version level as the IBM Spectrum Control server.

HWNCA0045E The plan with scheduleID *scheduleID* can not be deleted because the delivery unit creation is already processing state.

Explanation

You cannot delete or modify a delivery unit that is in a processing state, so its associated plan can not be deleted.

Action

Wait until the associated creation process is complete then try to delete the plan again.

HWNCA0046E The specified value *deletionType* is not a valid plan deletion type.

Explanation

You cannot delete a plan with a invalid deletion type.

Action

Pick a valid deletion type and then try to delete the plan again.

HWNCA0047E Host with host name *host* is not known by IBM Spectrum Control.

Explanation

No host with the specified name exists.

Action

HWNCA0048E Filer Storage ID *filerId* for given fileSystem is not found.

Explanation

Filer Id for a given fileSystem is either not found or not specified.

Action

HWNCA0049W IBM Spectrum Control did not change any Fibre Channel zones. After the provisioning operation completes, verify the Fibre Channel connectivity between any involved host and the storage subsystem that contains the storage volume. If necessary, change the zoning configuration.

Explanation

IBM Spectrum Control reassigned volumes to hosts and host ports, but did not change the fabric zoning to ensure the connectivity between the storage subsystem and the hosts.

Action

After the provisioning operation completes, configure the zones and zone sets, if necessary, using the appropriate external applications.

HWNCA0050W IBM Spectrum Control did not enable the multipath policy for the host or hosts to which the volume was reassigned. You must enable the multipath policy after the provisioning operation completes.

Explanation

IBM Spectrum Control reassigned volumes to hosts and host ports, but did not enable the multipath policy for the involved hosts.

Action

After the provisioning operation completes, log in to the operating system on each involved host to enable the multipath policy.

HWNCA0051E Incorrect volume name *volumeName*. The name can contain letters, numbers, dashes (-) and underscores (_). However, the name cannot start with a number, dash, or the reserved words *vdisk* or *volume*.

Explanation

The volume name is incorrect.

Action

Choose a valid volume name.

HWNCA0052E One or more of the requested volume names are too long for the following storage systems: *StorageSubsystemsList*. The maximum volume name length that is allowed is *MaxAllowedNameLength* characters, and a volume name of *LargestNameLength* characters was requested.

Explanation

Restrictions on volume name length vary by storage system type. When you request volumes, IBM Spectrum Control uses criteria such as the service class requirements, available space and performance data to identify the best location for the volumes. It does not consider the allowed length for volume names on particular storage systems. Sometimes a storage pool that is recommended for volume placement is on a storage system that does not support the volume name length.

Action

Either specify volume names within the allowed length, or select a capacity pool that contains only storage pools from storage subsystems that allow longer volume names.

HWNCA0053E The service class *scName* does not support provisioning from any available storage.

Explanation

This service class is restricted to a certain set of capacity pools, and therefore does not support provisioning from any available storage.

Action

Specify a capacity pool that can be used with this service class.

HWNCA0054E The volume *volumeName* already exists on the storage system *storageSystemName*. Request a different volume name.

Explanation

The best location for the requested volume already has a volume of that name. Because volume names must be unique on a storage system, the new volume cannot be created.

Action

Specify a different name for the volume.

HWNCA0055E The share *shareName* already exists on the storage system *storageSystemName*. Request a different share name.

Explanation

The best location for the requested NAS share already has a share of that name. Because share names must be unique on a storage system, the new share cannot be created.

Action

Specify a different name for the share.

HWNCA0056W Zoning cannot be configured because the connected fabrics are not monitored.

Explanation

Typically, you cannot provision volumes to servers and hypervisors that do not have fabric connectivity. However, because the selected hosts have Fibre Channel Port WWPNs and do not appear to have fabric connectivity, it is possible that the related fabrics were not probed by IBM Spectrum Control.

In this case, you can still provision the volumes, but because no fabric information is available, all fabric-related options are ignored and the fabric is not configured. Fabric-related options include number of paths, redundant fabrics, and automatic zoning.

Action

Ensure that the hosts are connected to the recommended storage system before you run the task. After the provisioning operation completes, you must manually configure the fabric.

HWNCA0057W The multipath policy cannot be set on the following hypervisors: *HostList*.

Explanation

IBM Spectrum Control cannot enable the multipath policy for hypervisors.

Action

After the provisioning operation completes, log in to the hypervisor or use VMware tools to enable the multipath policy on the hypervisor.

HWNCA0058W The multipath policy cannot be set on the following agentless servers: *HostList*.

Explanation

IBM Spectrum Control cannot enable the multipath policy for servers that are not managed by a Storage Resource agent.

Action

After the provisioning operation completes, log in to the operating system on the agentless server to enable the multipath policy.

HWNCA0059W The multipath policy cannot be set on the following hosts: *HostList*.

Explanation

IBM Spectrum Control cannot enable the multipath policy for hypervisors or servers that are not managed by a Storage Resource agent.

Action

For each hypervisor: After the provisioning operation completes, log in to the hypervisor or use VMware tools to enable the multipath policy on the hypervisor.

For each agentless server: After the provisioning operation completes, log in to the operating system on the agentless server to enable the multipath policy.

HWNCA0060E Storage volume can not be assigned to servers that are running different operating systems.

Explanation

A storage volume can be assigned to multiple servers only if the servers use the same operating systems.

Action

Ensure that the specified servers are running the same operating system, and try to create or modify the volume again.

HWNCA0061E The required resources cannot be reserved because too many provisioning tasks are currently being created. Try again in a few minutes.

Explanation

When a task for provisioning volumes is created, pool space is reserved for allocating the volumes. Because too many provisioning tasks are being created at the same time, pool space cannot be reserved and this provisioning task cannot be created.

Action

Wait a few minutes for the other provisioning tasks to be created.

HWNCA0062E The operation did not complete because the specified service class or capacity pool no longer exists.

Explanation

The service class or capacity pool that was specified in the provisioning request was deleted before IBM Spectrum Control could determine the recommended placement for a new volume or share. If either the service class or capacity pool no longer exists, the operation to create the provisioning task fails.

Action

Resubmit your provisioning request. Make sure that any service class or capacity pool that is specified in the provisioning request still exists.

HWNCA0063E The service class *ScName* that is specified by this provisioning task does not exist.

Explanation

The service class that was specified in the provisioning task was deleted before IBM Spectrum Control could complete the provisioning of the new volume or share. The provisioning task failed because the service class no longer exists.

Action

Recreate the provisioning task. Ensure that you specify an existing service class.

HWNCA0064E The current SMAC API client is version *current_version*. The requested operation requires SMAC API client version *required_version* or later.

Explanation

Some SMAC API operations require the most recent version of the SMAC client. This prevents unexpected behavior by older SMAC API clients.

Action

Update the SMAC API client to the version specified in this message, or to a more recent version. For best results, the SMAC API client should be at the same version level as the IBM Spectrum Control server.

HWNCA0065E The delivery unit specified has snapshots and could not be deleted.

Explanation

A delivery unit cannot be deleted if it has snapshots.

Action

Delete the delivery unit snapshots and try the deletion of the delivery unit again.

HWNCA0066E The WWPN *WWPN* already belongs to a server or hypervisor managed by IBM Spectrum Control. Was expected an unknown WWPN based on previous Hosts or WWPNS passed in.

Explanation

The WWPNS specified in the provisioning task was sent part of a group of WWPNS or Hosts that are not of same type. All specified WWPNS or Hosts should be either known or unknown. The provisioning task failed because the of this.

Action

Recreate the provisioning task. Ensure that you all wwpns and hosts of same type (known or unknown to IBM Spectrum Control). In case unknown wwpns are sent specify maximum one unknown host. If one unknown host is specified make sure at least one unknown wwpn is specified also.

HWNCA0067E The host *host* already belongs to a server or hypervisor managed by IBM Spectrum Control. Was expected an unknown Host based on previous Hosts or WWPNS passed in.

Explanation

The host specified in the provisioning task was sent part of a group of WWPNS or Hosts that are not of same type. All specified WWPNS or Hosts should be either known or unknown. The provisioning task failed because the of this.

Action

Recreate the provisioning task. Ensure that you all wwps and hosts of same type (known or unknown to IBM Spectrum Control). In case unknown wwps are sent specify maximum one unknown host. If one unknown host is specified make sure at least one unknown wwpn is specified also.

HWNCA0068E The list of hosts was expected to contain just one or no element because the WWPNS passed in were all unknown. The size of the host list is size.

Explanation

The host specified in the provisioning task was sent part of a group of WWPNS that are all Unknown to IBM Spectrum Control Server. Therefore only a maximum of one host (also unknown to IBM Spectrum Control) can be specified.

Action

Recreate the provisioning task. Ensure that you all wwps and hosts of same type (known or unknown to IBM Spectrum Control). In case unknown wwps are sent specify maximum one unknown host. If one unknown host is specified make sure at least one unknown wwpn is specified also.

HWNCA0069E No unknown WWPNS were specified for this provisioning operation but an unknown Host was specified.

Explanation

The unknown host specified in the provisioning task was not sent part of a group of unknown WWPNS.

Action

Recreate the provisioning task. Ensure that you all wwps and hosts of same type (known or unknown to IBM Spectrum Control). In case unknown wwps are sent specify maximum one unknown host. If one unknown host is specified make sure at least one unknown wwpn is specified also.

HWNCA0070E The expansion of the capacity is not supported for the delivery unit with ID *delivery_unit_id*.

Explanation

The storage subsystem containing the delivery unit does not support to expand the capacity.

Action

HWNCA0071E The reduction of the capacity is not supported for the delivery unit with ID *delivery_unit_id*.

Explanation

The storage subsystem containing the delivery unit does not support to reduce the capacity.

Action

HWNDA - Data Manager API messages

- [HWNDA0001I Operation Name of the operation processed successfully.](#)
- [HWNDA0002E Mandatory parameter Name of the mandatory parameter which is missing missing](#)
- [HWNDA0003E Invalid parameter Name of the parameter which was invalid](#)
- [HWNDA0004E An internal error occurred.](#)
- [HWNDA0005E The server encountered an error when it was accessing the database.](#)
- [HWNDA0006E The name provided while creating a new group is already in use.](#)
- [HWNDA0007E An external key could not be identified for the provided type The constant integer type of the Group element and id The unique integer database ID of the Group element.](#)

- [HWNDA0008E The specified attribute invalid attribute name is not a valid attribute.](#)
- [HWNDA0009E An internal ID could not be identified for the provided type The constant integer type of the Group element and external key The unique external key of the Group element.](#)
- [HWNDA0010I The following elements are already members of the group The group: The keys of the elements.](#)
- [HWNDA0011I The following elements are not members of the The group Group and cannot be removed: The element key.](#)
- [HWNDA0012E Adding a Group with the name Name of the proposed new member to the Name of the parent group Group would create a circular relationship that is not allowed.](#)
- [HWNDA0013E The input parameter value input parameter value for input input parameter name exceeds the maximum allowable length of number of allowable characters characters.](#)
- [HWNDA0014E The provided Group attribute value Group attribute value for the Group attribute name Group attribute contains invalid characters. The following characters are not allowed, \\/:*?><|,."](#)
- [HWNDA0015E You are not the original creator of the provided Group name Group name.](#)
- [HWNDA0016E The provided Tiering Policy name Tiering Policy name is not unique.](#)
- [HWNDA0017E The provided Group Group name or ID does not exist.](#)
- [HWNDA0018E The provided Tiering Policy name Tiering Policy name does not exist.](#)
- [HWNDA0019E The provided candidate and destination Group names, Group name, cannot be the same.](#)
- [HWNDA0020E The provided condition condition type is not valid.](#)
- [HWNDA0021E The provided operand operand type is not valid.](#)
- [HWNDA0022E The provided condition condition type is either already applied to this tiering policy or conflicts with an existing condition, existing condition type.](#)
- [HWNDA0023E The requested priority value priority value is invalid.](#)
- [HWNDA0024E The specified Group name Group name is not unique.](#)
- [HWNDA0025E Cannot add the specified resource because the resource type, element type, is not supported as a child of the group.](#)
- [HWNDA0026E Cannot add the specified group, Group name, because the group type, type, is not supported as a child of the application.](#)
- [HWNDA0027E The first option specified in the file must be -appgroupname.](#)
- [HWNDA0028E The argument of the option option is missing at or before line Line Number : Line](#)
- [HWNDA0029E Both option1 and option2 were specified at or before line Line Number : Line](#)
- [HWNDA0030E The option option is missing at or before line Line Number : Line](#)
- [HWNDA0031E Neither option1 nor option2 was specified at or before line Line Number : Line](#)
- [HWNDA0032E Invalid number of parameters for option option at or before line Line Number : Line](#)
- [HWNDA0033E Incomplete options sequence before end of file.](#)
- [HWNDA0034E Invalid option option at line Line Number : Line](#)
- [HWNDA0035E Invalid resource type type at line Line Number : Line](#)
- [HWNDA0036E Invalid sequence of options at or before line Line Number : Line](#)
- [HWNDA0037E Syntax error, quote sequence not properly closed at line Line Number : Line](#)
- [HWNDA0038E Option option is not allowed for resource type type at or before line Line Number : Line](#)
- [HWNDA0039E The input data for modifying the application groups is missing.](#)
- [HWNDA0040E An invalid element was encountered in the input data.](#)
- [HWNDA0041E The application group name is missing from the input data.](#)
- [HWNDA0042E The operation is missing from the input data.](#)
- [HWNDA0043E The resource type is missing from the input data.](#)
- [HWNDA0044E The server name is missing from the input data.](#)
- [HWNDA0045E The device name is missing from the input data.](#)
- [HWNDA0046E Invalid values were specified for the server, device or cluster names in the input data.](#)
- [HWNDA0047E Member names and the tags were specified for the same operation.](#)
- [HWNDA0048W The following entities were not found: Entities.](#)
- [HWNDA0049W No entities were found for the tags: Tags](#)
- [HWNDA0050E The member names or tags were not specified for the operation.](#)
- [HWNDA0051W The group Name of the proposed new member cannot be added to itself.](#)
- [HWNDA0052W The group Name of the proposed new member cannot be added to the Name of the parent group group because it creates a circular relationship that is not allowed.](#)
- [HWNDA0053W The group Name of the group contains child groups and cannot be deleted. The child groups must be deleted before the parent group can be deleted.](#)
- [HWNDA0054E The filter mask that was used to create or edit a group filter is currently being used.](#)
- [HWNDA0055E The argument argument for the parameter parameter on line Line Number : Line is invalid.](#)
- [HWNDA0056E The first option specified in the file must be -id.](#)
- [HWNDA0057E The specified Group Group name is not of type type.](#)
- [HWNDA0058W These groups have same names as existing members of the group The group: The keys of the elements. They were not added to the group.](#)
- [HWNDA0059E The specified tag key Tag key is not valid.](#)
- [HWNDA0060W These group members are also members of another group: The groups. They are not deleted.](#)
- [HWNDA0061W These group members cannot be moved up one level in hierarchy due to name conflicts: The groups. The group is not deleted.](#)
- [HWNDA0062E The specified Group Group name is not a department group.](#)
- [HWNDA0063W These group members cannot be moved as top level groups in hierarchy due to name conflicts: The groups. No group members were removed from the group.](#)
- [HWNDA0064E An application with the same name already exists.](#)
- [HWNDA0065E A department with the same name already exists.](#)
- [HWNDA0066E Invalid values were specified for the device, cluster or file system names in the input data.](#)
- [HWNDA0067E Resources of type Type of resource cannot be added to an application or removed from an application using tags.](#)
- [HWNDA0068E The application cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.](#)
- [HWNDA0069E The department cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.](#)
- [HWNDA0070W The File Systems: {0} were not added to the application {1} because they are NAS file systems.](#)
- [HWNDA0071E Member IDs and the tags were specified for the same operation.](#)
- [HWNDA0072E Member IDs should be specified in the input data for resources of type appgroup.](#)
- [HWNDA0073E A general group with the same name already exists.](#)
- [HWNDA0074E The general group cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.](#)
- [HWNDA0075E A dashboard with the same name already exists.](#)
- [HWNDA0076E The dashboard group cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.](#)
- [HWNDA0077E A policy group with the same name already exists.](#)
- [HWNDA0078I The Name of the policy group policy group was removed.](#)

HWNDA0001I *Operation Name of the operation processed successfully.*

Explanation

A request was run successfully. No error condition was encountered.

Action

No action is required.

HWNDA0002E *Mandatory parameter Name of the mandatory parameter which is missing missing*

Explanation

Mandatory parameter {0} is missing.

Action

Contact IBM Support.

Related reference

- [Getting support](#)

HWNDA0003E *Invalid parameter Name of the parameter which was invalid*

Explanation

Invalid parameter {0} was encountered while processing an API request.

Action

Contact IBM Support.

Related reference

- [Getting support](#)

HWNDA0004E *An internal error occurred.*

Explanation

An internal error occurred during execution.

Action

Check the logs for an indication of an error or exception and contact IBM customer support.

Related reference

- [Getting support](#)

HWNDA0005E *The server encountered an error when it was accessing the database.*

Explanation

The server cannot access data that is stored in the database.

Action

Check the status of the database. Also, check the logs for an indication of an error or exception and contact IBM customer support.

Related reference

- [Getting support](#)

HWNDA0006E The name provided while creating a new group is already in use.

Explanation

The name of a new group must be unique.

Action

Choose a unique name for the group and try again.

HWNDA0007E An external key could not be identified for the provided type *The constant integer type of the Group element* and id *The unique integer database ID of the Group element*.

Explanation

The element that is a member of a Group could not be found in the database and indicates some form of database inconsistency.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

HWNDA0008E The specified attribute *invalid attribute name* is not a valid attribute.

Explanation

The attribute that is specified is not in the list of usable attributes for the method that is being invoked.

Action

Contact IBM customer support.

Related reference

- [Getting support](#)

HWNDA0009E An internal ID could not be identified for the provided type *The constant integer type of the Group element* and external key *The unique external key of the Group element*.

Explanation

The provided external key could not be translated into an internal database ID.

Action

Verify the validity of the external key.

HWNDA0010I The following elements are already members of the group
The group: The keys of the elements.

Explanation

The groups cannot contain the same member more than once.

Action

Verify the current contents of the group.

HWNDA0011I The following elements are not members of the *The group*
Group and cannot be removed: *The element key.*

Explanation

The element that was requested to be removed is not currently a member of the Group.

Action

Verify that the current contents of the Group are correct.

HWNDA0012E Adding a Group with the name *Name of the proposed new member* to the *Name of the parent group* Group would create a circular relationship that is not allowed.

Explanation

A Group may not be a member of its children.

Action

Verify that the Group is not being added to any of its children groups.

HWNDA0013E The input parameter value *input parameter value* for
input *input parameter name* exceeds the maximum allowable length of
number of allowable characters characters.

Explanation

The provided input is too long.

Action

Shorten the provided input and try again.

HWNDA0014E The provided Group attribute value *Group attribute value* for the *Group attribute name* Group attribute contains

invalid characters. The following characters are not allowed, \\/:*?><|."

Explanation

The proposed Group attribute contained an invalid character.

Action

Remove the invalid character or characters and try the action again.

HWNDA0015E You are not the original creator of the provided Group name *Group name*.

Explanation

You must be the original creator of the Group in order to update its properties.

Action

None.

HWNDA0016E The provided Tiering Policy name *Tiering Policy name* is not unique.

Explanation

The proposed Tiering Policy name is already in use.

Action

Either specify a new name or delete the current Tiering Policy with this name and rerun this command.

HWNDA0017E The provided Group *Group name or ID* does not exist.

Explanation

The Data server could not find a Group with the specified name or ID.

Action

Verify that the required Group exists. If it does not, specify a different Group.

HWNDA0018E The provided Tiering Policy name *Tiering Policy name* does not exist.

Explanation

The Data server could not find a Tiering Policy with the name specified.

Action

Verify that the required Tiering Policy exists. If it does not, create a Tiering Policy with the required name or specify a different Tiering Policy.

HWNDA0019E The provided candidate and destination Group names, *Group name*, cannot be the same.

Explanation

The Tiering Policy cannot have the same Group for both its candidate and destination.

Action

Choose a different Group for either the candidate or destination for this Tiering Policy.

HWNDA0020E The provided condition *condition* type is not valid.

Explanation

Only valid condition types may be specified when adding a new condition to a tiering policy.

Action

Specify a valid condition type to add to the tiering policy.

HWNDA0021E The provided operand *operand type* is not valid.

Explanation

Only valid operands may be specified when adding a new condition to a tiering policy.

Action

Specify a valid operand for the condition being added to the tiering policy.

HWNDA0022E The provided condition *condition type* is either already applied to this tiering policy or conflicts with an existing condition, *existing condition type*.

Explanation

Tiering policies may only contain one condition. The exception is that an AGE and an IODENSITY conditions may exist for the same tiering policy.

Action

Remove the existing condition from the tiering policy before adding the new condition.

HWNDA0023E The requested priority value *priority value* is invalid.

Explanation

Tiering policy priorities range from 1 (highest priority) to the current number of policies.

Action

Try the command again, specifying a valid priority value.

HWNDA0024E The specified Group name *Group name* is not unique.

Explanation

The specified Group name is not unique and none of these Groups are owned by the current user.

Action

Try the command again, specifying the full Group name in "user"."name" format.

HWNDA0025E Cannot add the specified resource because the resource type, *element type*, is not supported as a child of the group.

Explanation

The specified resource type is not supported as a child of the group.

Action

Check the resource types that you are allowed to add as a child for this group. Try the operation again, using a supported resource type.

HWNDA0026E Cannot add the specified group, *Group name*, because the group type, *type*, is not supported as a child of the application.

Explanation

The specified group is not supported as a child of an application.

Action

Check the group types that you are allowed to add as a child of an application. Try the operation again, using a supported group type.

HWNDA0027E The first option specified in the file must be -*appgroupname*.

Explanation

The first option specified in the file must be -*appgroupname*.

Action

Edit the file appropriately.

HWNDA0028E The argument of the option *option* is missing at or before line *Line Number* : *Line*

Explanation

The option argument was not specified.

Action

Edit the file appropriately.

HWNDA0029E Both *option1* and *option2* were specified at or before line *Line Number* : *Line*

Explanation

Both options were specified.

Action

Edit the file appropriately.

HWNDA0030E The option *option* is missing at or before line *Line Number : Line*

Explanation

An option was not specified.

Action

Edit the file appropriately.

HWNDA0031E Neither *option1* nor *option2* was specified at or before line *Line Number : Line*

Explanation

Action

Edit the file appropriately.

HWNDA0032E Invalid number of parameters for option *option* at or before line *Line Number : Line*

Explanation

An Invalid number of parameters were specified for an option.

Action

Edit the file appropriately.

HWNDA0033E Incomplete options sequence before end of file.

Explanation

An incomplete options sequence was detected.

Action

Edit the file appropriately.

HWNDA0034E Invalid option *option* at line *Line Number : Line*

Explanation

An Invalid option was specified.

Action

Edit the file appropriately.

HWNDA0035E Invalid resource type *type* at line *Line Number* : *Line*

Explanation

An invalid resource type was specified.

Action

Edit the file appropriately.

HWNDA0036E Invalid sequence of options at or before line *Line Number* : *Line*

Explanation

An invalid sequence of options was specified.

Action

Edit the file appropriately.

HWNDA0037E Syntax error, quote sequence not properly closed at line *Line Number* : *Line*

Explanation

The quote sequence was not properly closed.

Action

Edit the file appropriately.

HWNDA0038E Option *option* is not allowed for resource type *type* at or before line *Line Number* : *Line*

Explanation

The option and the resource type are not compatible.

Action

Edit the file appropriately.

HWNDA0039E The input data for modifying the application groups is missing.

Explanation

No input data was specified to modify the application groups.

Action

Specify the necessary input data.

HWNDA0040E An invalid element was encountered in the input data.

Explanation

An invalid element was encountered while processing an API request for modifying the application groups.

Action

Specify a correct element.

HWNDA0041E The application group name is missing from the input data.

Explanation

An application group name was not specified.

Action

Specify a correct application group name.

HWNDA0042E The operation is missing from the input data.

Explanation

An operation was not specified.

Action

Specify a correct operation.

HWNDA0043E The resource type is missing from the input data.

Explanation

A resource type was not specified.

Action

Specify a correct resource type.

HWNDA0044E The server name is missing from the input data.

Explanation

A server name was missing while processing an API for modifying the application groups.

Action

Specify a correct server name.

HWNDA0045E The device name is missing from the input data.

Explanation

A device name was missing while processing an API for modifying the application groups.

Action

Specify a correct device name.

HWNDA0046E Invalid values were specified for the server, device or cluster names in the input data.

Explanation

Invalid values were specified for the server, device or cluster names in the input data that is used to modify the applications.

Action

Specify valid values for the server, device and cluster names.

HWNDA0047E Member names and the tags were specified for the same operation.

Explanation

Both the member names and tags were specified for the same operation.

Action

Specify either member names or tags for an operation.

HWNDA0048W The following entities were not found: *Entities*.

Explanation

The entities were not found in the database.

Action

Try the command again and specify the correct names for the entities.

HWNDA0049W No entities were found for the tags: *Tags*

Explanation

The entities were not found in the database.

Action

Try the command again and specify the correct tag keys and values.

HWNDA0050E The member names or tags were not specified for the operation.

Explanation

The members need to be specified using the individual names or the tag pairs.

Action

Specify either member names or tags for an operation.

HWNDA0051W The group *Name of the proposed new member* cannot be added to itself.

Explanation

You cannot add a group to itself. Your request to add a group to itself is ignored.

Action

Verify the user defined properties of the group.

HWNDA0052W The group *Name of the proposed new member* cannot be added to the *Name of the parent group* group because it creates a circular relationship that is not allowed.

Explanation

Groups cannot be used in a circular relationship with their parent and child groups.

Action

Verify the current contents of the group are correct.

HWNDA0053W The group *Name of the group* contains child groups and cannot be deleted. The child groups must be deleted before the parent group can be deleted.

Explanation

A Storage Resource Group, Reporting Group or Application Group that has child groups cannot be deleted.

Action

Delete the child groups before you delete any parent group.

HWNDA0054E The filter mask that was used to create or edit a group filter is currently being used.

Explanation

The filter mask must be unique for each group.

Action

Use a unique filter mask for the group filter and try to create or edit the group again.

HWNDA0055E The argument *argument* for the parameter *parameter* on line *Line Number* : *Line* is invalid.

Explanation

The argument for the parameter is not valid.

Action

Refer to the command help for valid arguments for the parameter and edit the input file to specify a valid argument.

HWNDA0056E The first option specified in the file must be *-id*.

Explanation

The first option specified in the file must be -id.

Action

Edit the file appropriately.

HWNDA0057E The specified Group *Group name* is not of type *type*.

Explanation

The specified Group is not of the specified type.

Action

Try the command again, specifying the correct type.

HWNDA0058W These groups have same names as existing members of the group *The group: The keys of the elements*. They were not added to the group.

Explanation

The group members must have unique names in order to be added to the group.

Action

Rename the group member you are trying to add or remove the group member with the same name that already exists.

HWNDA0059E The specified tag key *Tag key* is not valid.

Explanation

The tag key that you specified does not exist.

Action

Specify a valid tag key or use the -key parameter when you modify an application group, or the -memberid parameter when you modify a department group.

HWNDA0060W These group members are also members of another group: *The groups*. They are not deleted.

Explanation

Even though you specified the -rmchildren parameter, the group members are not deleted because they also belong to another group.

Action

Remove the group member from each group that it belongs to using the modifydeptgroup command.

HWNDA0061W These group members cannot be moved up one level in hierarchy due to name conflicts: *The groups*. The group is not deleted.

Explanation

When you do not specify the `--rmchildren` parameter, the group members are moved up one level in hierarchy. The specified group members have the same name as the existing groups in the higher hierarchical level.

Action

Rename the specified groups and try the command again.

HWNDA0062E The specified Group *Group name* is not a department group.

Explanation

The specified group is not of type department.

Action

Specify a valid department group and try the command again.

HWNDA0063W These group members cannot be moved as top level groups in hierarchy due to name conflicts: *The groups*. No group members were removed from the group.

Explanation

When you remove group members from a group of the same type, and that group is their only single parent, the group members must become the top level groups in the hierarchy. The specified group members have the same name as the existing top level groups in hierarchy.

Action

Rename the specified groups and try the command again.

HWNDA0064E An application with the same name already exists.

Explanation

The specified application name already exists and cannot be duplicated.

Action

Enter a unique name for the application.

HWNDA0065E A department with the same name already exists.

Explanation

The specified department name already exists and cannot be duplicated.

Action

Enter a unique name for the department.

HWNDA0066E Invalid values were specified for the device, cluster or file system names in the input data.

Explanation

Invalid values were specified for the device, cluster or file system names in the input data that is used to modify the applications.

Action

Specify valid values for the device, cluster and file system names.

HWNDA0067E Resources of type *Type of resource* cannot be added to an application or removed from an application using tags.

Explanation

Resources of this type can only be added to an application or removed from an application by specifying the resource key.

Action

Retry adding or removing the resource by specifying the resource key.

HWNDA0068E The application cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.

Explanation

When you added a new subcomponent for the application, the maximum number of levels that are allowed in the hierarchy was exceeded.

Action

Review the application hierarchy. Add the application as a child of an application that is at a higher level in the hierarchy. Alternatively, add the application as a top-level application.

HWNDA0069E The department cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.

Explanation

When you created a new subdepartment for the department, the maximum number of levels that are allowed in the hierarchy was exceeded.

Action

Review the department hierarchy. Add the department as a child of a department that is at a higher level in the hierarchy. Alternatively, add the department as a top-level department.

HWNDA0070W The File Systems: {0} were not added to the application {1} because they are NAS file systems.

Explanation

A NAS file system is not a supported type to add to an application. Only server file systems are supported.

Action

Add a supported type of file system to the application.

HWNDA0071E Member IDs and the tags were specified for the same operation.

Explanation

Both the member IDs and tags were specified for the same operation.

Action

Specify either member IDs or tags for an operation.

HWNDA0072E Member IDs should be specified in the input data for resources of type appgroup.

Explanation

Member IDs of the resources of type appgroup were not specified in the input data that is used to modify the applications.

Action

Specify valid member IDs.

HWNDA0073E A general group with the same name already exists.

Explanation

The specified general group name already exists and cannot be duplicated.

Action

Enter a unique name for the general group.

HWNDA0074E The general group cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.

Explanation

When you added a new subcomponent for the general group, the maximum number of levels that are allowed in the hierarchy was exceeded.

Action

Review the general group hierarchy. Add the general group as a child of a group that is at a higher level in the hierarchy. Alternatively, add the general group as a top-level group.

HWNDA0075E A dashboard with the same name already exists.

Explanation

The specified dashboard name already exists and cannot be duplicated.

Action

Enter a unique name for the dashboard.

HWNDA0076E The dashboard group cannot be added because the number of levels in the group hierarchy exceeds the maximum of five levels.

Explanation

When you added a new subcomponent for the dashboard group, the maximum number of levels that are allowed in the hierarchy was exceeded.

Action

Review the dashboard group hierarchy. Add the dashboard group as a child of a group that is at a higher level in the hierarchy. Alternatively, add the dashboard group as a top-level group.

HWNDA0077E A policy group with the same name already exists.

Explanation

The specified policy group name already exists and cannot be duplicated.

Action

Enter a unique name for the policy group.

HWNDA0078I The *Name of the policy group* policy group was removed.

Explanation

This message is for informational purposes only.

Action

No further action is required.

HWNEM - Element manager messages

- [HWNEM0001E The element manager management service failed to obtain a database connection.](#)
- [HWNEM0002E An error occurred while attempting to add element manager information to the database.](#)
- [HWNEM0003E An error occurred while attempting to update element manager information in the database.](#)
- [HWNEM0004E An error occurred while attempting to remove element manager information from the database.](#)
- [HWNEM0005E An error occurred while attempting to obtain element manager information from the database.](#)
- [HWNEM0006E Failed to transmit request to Data Server to initiate SMI-S provider discovery job.](#)
- [HWNEM0007E An error occurred while attempting to obtain Data Server information from the database.](#)
- [HWNEM0008E Data Server information was not found in the database. SMI-S provider discovery could not be scheduled.](#)
- [HWNEM0009E An error occurred while attempting to obtain the SMI-S provider URL associated with an element manager from the database.](#)
- [HWNEM0010E An error occurred while attempting to obtain element manager credentials from the database.](#)
- [HWNEM0011E An error occurred while attempting to remove element manager credentials from the database.](#)
- [HWNEM0012E An error occurred while attempting to store element manager credentials in the database.](#)
- [HWNEM0013E Encountered element manager with malformed URL \(URL\).](#)
- [HWNEM0014E Encountered element manager with URL containing hostname that could not be resolved by DNS \(URL\).](#)
- [HWNEM0015E Failed to clone element manager.](#)
- [HWNEM0016E Failed to encrypt element manager password.](#)
- [HWNEM0017E Failed to decrypt element manager password.](#)
- [HWNEM0018E Failed to send request to Data Server \(dataServerHost:dataServerPort\).](#)
- [HWNEM0019I Attempting to schedule discovery on Data Server for number SMI-S provider\(s\) \(SMI-S provider URLs\)...](#)
- [HWNEM0020I Successfully scheduled discovery on Data Server for number SMI-S provider\(s\).](#)
- [HWNEM0021I SMI-S provider discovery was not scheduled on Data Server. No SMI-S providers are associated with the specified set of element managers.](#)
- [HWNEM0022E Failed to authenticate with ESSNI server associated with element manager at URL using ESSNI user ID ESSNI user ID.](#)
- [HWNEM0023E The ESSNI server associated with the element manager at URL is not available.](#)
- [HWNEM0026E An error occurred while attempting to retrieve version information from the ESSNI server associated with the element manager at URL.](#)
- [HWNEM0029E Failed to set element manager credentials. Specified Element manager does not exist.](#)
- [HWNEM0100E The element manager's URL is not in the correct format.](#)
- [HWNEM0101E Change the default administrator password now to avoid security conflicts.](#)
- [HWNEM0102E The element manager already exists.](#)
- [HWNEM0103E A problem occurred adding the element manager to IBM Spectrum Control.](#)
- [HWNEM0104E Login to the element manager failed. Check the user credentials.](#)
- [HWNEM0105E The request contained data in an unexpected format](#)
- [HWNEM0106E The request did not contain the expected information.](#)
- [HWNEM0107E The request contained an unsupported action.](#)
- [HWNEM0108E Connection test to SMI-S provider URL FAILED due to status code.](#)
- [HWNEM0109E Connection test to Element Manager failed.](#)
- [HWNEM0111E A problem occurred locating the element manager in IBM Spectrum Control.](#)
- [HWNEM0112E A problem occurred updating the element manager to IBM Spectrum Control.](#)
- [HWNEM0113E A problem occurred removing the SMI-S provider from the element manager.](#)
- [HWNEM0114E A problem occurred locating the SMI-S provider associated with the element manager.](#)
- [HWNEM0115E Unable to establish an https connection to the element manager.](#)
- [HWNEM0116E A problem occurred removing the element manager from IBM Spectrum Control.](#)
- [HWNEM0117E A problem occurred testing the connection to the element manager.](#)

- [HWNEM0118I Connection test to the element manager element manager passed.](#)
- [HWNEM0119E Connection test to the element manager element manager failed.](#)
- [HWNEM0120E IBM Spectrum Control is unable to communicate with the element manager, the problem could be the element manager is not running, a network communication error or user credentials stored in IBM Spectrum Control are incorrect.](#)
- [HWNEM0121I Connection test to the SMI-S provider SMI-S provider passed.](#)
- [HWNEM0122I Connection test to the SMI-S provider SMI-S provider failed.](#)
- [HWNEM0123E Before executing the action the user must select an element manager from the table.](#)
- [HWNEM0124E There is no SMI-S provider associated with the selected element manager.](#)
- [HWNEM0125I The element manager is about to be removed. Once removed the element manager will not be accessible from IBM Spectrum Control. Do you wish to continue.](#)
- [HWNEM0126I The element manager is about to be removed. Once removed the element manager will not be accessible from IBM Spectrum Control. Do you wish to continue.](#)
- [HWNEM0127E An internal processing error occurred while servicing the last request.](#)
- [HWNEM0128E The element manager is not available. Ensure that the element manager's credentials are defined and up to date.](#)
- [HWNEM0129E An unexpected error occurred changing the element manager's default password.](#)
- [HWNEM0130E SMI-S provider connection was added to the element manager, however the SMI-S provider discovery job failed to launch.](#)
- [HWNEM0131I The DS8000 SMI-S provider has been added successfully, IBM Spectrum Control has started the discovery job for the DS8000 storage subsystem managed by this SMI-S provider. To check the status of the jobs, go to the IBM Spectrum Control perspective and check the following navigation tree nodes: Administrative Services -> Discovery -> CIMOM](#)
- [HWNEM0132I SMI-S provider connection was removed from element manager.](#)
- [HWNEM0133E Internal error occurred, the element manager info for launching the element manager could not be located in the IBM Spectrum Control DB.](#)
- [HWNEM0134E The user session data is no longer available. To continue restart the IBM Spectrum Control GUI.](#)
- [HWNEM0135E Unable to contact the device server. It appears to be down.](#)
- [HWNEM0136I The DS8000 Element Manager at IP Address has been added successfully.](#)
- [HWNEM0137I service.method IBM Spectrum Control User IBM Spectrum Control user launched DS8000 Element Manager IP Address under the alias of DS8000 user name.](#)
- [HWNEM0138I The DS8000 SMI-S provider has been modified successfully, IBM Spectrum Control has started the discovery job for the DS8000 storage subsystem managed by this SMI-S provider. To check the status of the jobs, go to the IBM Spectrum Control perspective and check the following navigation tree nodes: Administrative Services -> Data Sources -> Discovery -> CIMOM](#)
- [HWNEM0139I The DS8000 SMI-S provider has been modified successfully.](#)
- [HWNEM0140W The element manager GUI cannot be accessed because no username and password have been specified by the currently logged-in IBM Spectrum Control user. Element manager credentials are managed on a per-IBM Spectrum Control user basis. The element manager GUI will appear after you provide the correct username and password information.](#)
- [HWNEM0141E The length of the SMI-S provider's text description is too long. The description should be 255 characters or less.](#)
- [HWNEM0142E An element manager named manager.name already exists.](#)
- [HWNEM0143E The DS8000 Element Manager Console is only accessible from IBM Spectrum Control. To access it, open IBM Spectrum Control and switch to the DS8000 Element Manager perspective.](#)

HWNEM0001E The element manager management service failed to obtain a database connection.

Explanation

The element manager management service failed to obtain a database connection.

Action

Verify that the database is up and running. Check the Device Server logs for error messages that might help figure out what is going on. If such has not already been done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0002E An error occurred while attempting to add element manager information to the database.

Explanation

A SQLException was thrown while attempting to insert element manager information into the database. The exception was caught. In response, the transaction was terminated and rolled back. Element manager information was not successfully added to the database.

Action

Verify that the database is up and running. Check the Device Server logs for error messages that might help figure out what is going on. If such has not already been done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0003E An error occurred while attempting to update element manager information in the database.

Explanation

A SQLException was thrown while attempting to update element manager information in the database. The exception was caught. In response, the transaction was terminated and rolled back. Element manager information was not successfully updated in the database.

Action

Verify that the database is up and running. Check the Device Server logs for error messages that might help figure out what is going on. If such has not already been done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0004E An error occurred while attempting to remove element manager information from the database.

Explanation

A SQLException was thrown while attempting to remove element manager information from the database. The exception was caught. In response, the transaction was terminated and rolled back. Element manager information was not successfully removed from the database.

Action

Verify that the database is up and running. Check the Device Server logs for error messages that might help figure out what is going on. If such has not already been done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0005E An error occurred while attempting to obtain element manager information from the database.

Explanation

A SQLException was thrown while attempting to obtain element manager information from the database. The exception was caught. Element manager information was not successfully retrieved from the database.

Action

Verify that the database is up and running. Check the Device Server logs for error messages that might help figure out what is going on. If such has not already been done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0006E Failed to transmit request to Data Server to initiate SMI-S provider discovery job.

Explanation

An error occurred while attempting to kick off an SMI-S provider discovery job on the Data Server. The request was not successful. An SMI-S provider discovery job was neither scheduled nor initiated on the Data Server.

Action

Verify that the Data Server is up and running. Check the both the Data Server and Device Server log files for error messages that might help figure out what is going on. If such has not already done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0007E An error occurred while attempting to obtain Data Server information from the database.

Explanation

A SQLException was thrown while attempting to obtain Data Server information from the database. The exception was caught. Data Server information was not successfully retrieved from the database.

Action

Verify that the database is up and running. Check the Device Server logs for error messages that might help figure out what is going on. If such has not already been done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0008E Data Server information was not found in the database. SMI-S provider discovery could not be scheduled.

Explanation

An attempt was made to retrieve Data Server information from the database. However, no such information was found. Communication with the Data Server is required to schedule SMI-S provider discovery. Given that information about a Data Server was not available, SMI-S provider discovery was not initiated successfully.

Action

Data Server information is stored in T_RES_SERVER. Check that table for a row in which the value in the SERVER_TYPE column is set to 0. If no such row exists, restart the Data Server. Doing so should result in the creation of an appropriate entry in T_RES_SERVER. If, however, the problem persists, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0009E An error occurred while attempting to obtain the SMI-S provider URL associated with an element manager from the database.

Explanation

An SQLException was thrown while attempting to obtain the service URL of the SMI-S provider associated with an element manager. The exception was caught. The SMI-S provider service URL was not successfully retrieved from the database.

Action

Verify that the database is up and running. Check the Device Server logs for error messages that might help figure out what is going on. If such has not already been done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0010E An error occurred while attempting to obtain element manager credentials from the database.

Explanation

An SQLException was thrown while attempting to obtain the set of credentials associated with an element manager. The exception was caught. The element manager credentials were not successfully retrieved from the database.

Action

Verify that the database is up and running. Check the Device Server logs for error messages that might help figure out what is going on. If such has not already been done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0011E An error occurred while attempting to remove element manager credentials from the database.

Explanation

An SQLException was thrown while attempting to remove element manager credentials from the database. The exception was caught. In response, the transaction was terminated and rolled back. The element manager credentials were not successfully removed from the database.

Action

Verify that the database is up and running. Check the Device Server logs for error messages that might help figure out what is going on. If such has not already been done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0012E An error occurred while attempting to store element manager credentials in the database.

Explanation

An SQLException was thrown while attempting to store element manager credentials in the database. The exception was caught. In response, the transaction was terminated and rolled back. The element manager credentials were not successfully stored in the database.

Action

Verify that the database is up and running. Check the Device Server logs for error messages that might help figure out what is going on. If such has not already been done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0013E Encountered element manager with malformed URL (URL).

Explanation

The URL associated with an element manager was invalid. Processing involving the element manager was not able to proceed.

Action

Look for the URL specified in the error message in the URL column of the T_RES_REGISTERED_EM table in the database. Attempt to update the URL so that it is no longer invalid. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0014E Encountered element manager with URL containing hostname that could not be resolved by DNS (URL).

Explanation

The hostname associated with the element manager could not be resolved in DNS. Processing involving the element manager was not able to proceed.

Action

Verify that the hostname specified in the error message looks correct and that it can be pinged from the Device Server machine. If it cannot be reached from the Device Server machine, check the DNS configuration of the Device Server machine and its surrounding environment. Should everything look good, and the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0015E Failed to clone element manager.

Explanation

An error occurred while attempting to clone an element manager.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0016E Failed to encrypt element manager password.

Explanation

An error occurred while attempting to encrypt an element manager password.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0017E Failed to decrypt element manager password.

Explanation

An error occurred while attempting to decrypt an element manager password.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

HWNEM0018E Failed to send request to Data Server (*dataServerHost:dataServerPort*) .

Explanation

An error occurred while attempting to send a request to the Data Server.

Action

Verify that the Data Server is up and running. Check the both the Data Server and Device Server log files for error messages that might help figure out what is going on. If such has not already done, consider turning on the maximum level of tracing to aid in resolution of the issue. Should the problem persist, contact IBM support.

Related reference

-  [Getting support](#)

HWNEM0019I Attempting to schedule discovery on Data Server for *number* SMI-S provider(s) (*SMI-S provider URLs*) ...

Explanation

Informational message.

Action

N/A.

HWNEM0020I Successfully scheduled discovery on Data Server for *number* SMI-S provider(s) .

Explanation

Informational message.

Action

N/A.

HWNEM0021I SMI-S provider discovery was not scheduled on Data Server. No SMI-S providers are associated with the specified set of element managers.

Explanation

Informational message.

Action

N/A.

HWNEM0022E Failed to authenticate with ESSNI server associated with element manager at *URL* using ESSNI user ID *ESSNI user ID*.

Explanation

The Device Server was not able to login to the ESSNI server associated with the element manager at the specified URL. The ESSNI user ID and / or password was / were incorrect.

Action

Update the credentials used to connect to the ESSNI server and try again. Should valid credentials be specified and the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0023E The ESSNI server associated with the element manager at *URL* is not available.

Explanation

The Device Server was not able to connect to the ESSNI server associated with the element manager at the URL specified in the message.

Action

Verify that the ESSNI server is up and running. Use a tool like 'ping' to verify that it can be reached from the Device Server. Should everything check out fine and the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0026E An error occurred while attempting to retrieve version information from the ESSNI server associated with the element manager at *URL*.

Explanation

An error occurred while attempting to retrieve version information from the ESSNI server associated with the element manager at the specified URL. The operation failed.

Action

Check the ESSNI server logs for error messages that might help figure out what is going on. Should the problem persist, contact IBM support.

Related reference

- [Getting support](#)

HWNEM0029E Failed to set element manager credentials. Specified Element manager does not exist.

Explanation

An attempt was made to set credentials for an element manager that does not exist. The attempt failed.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0100E The element manager's URL is not in the correct format.

Explanation

The element manager could not be added to the management console because the information used to construct the URL of the element manager's administration console was wrong.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0101E Change the default administrator password now to avoid security conflicts.

Explanation

Retaining a password with its known default value is an insecure practice.

Action

Change the password to a value other than the default.

HWNEM0102E The element manager already exists.

Explanation

The user is attempting to add a duplicate of an element manager already defined in IBM Spectrum Control.

Action

Change the settings so that the element manager you are trying to add is unique.

HWNEM0103E A problem occurred adding the element manager to IBM Spectrum Control.

Explanation

Some unexpected error occurred when attempting to add the element manager.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0104E Login to the element manager failed. Check the user credentials.

Explanation

The element manager's credentials cannot be authenticated.

Action

Make sure you are entering the element manager username and password correctly.

HWNEM0105E The request contained data in an unexpected format

Explanation

The servlet received a request with data in an unexpected format.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0106E The request did not contain the expected information.

Explanation

The servlet received a request that did not contain the expected data.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0107E The request contained an unsupported action.

Explanation

The servlet received a request containing an undefined action.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0108E Connection test to SMI-S provider *URL FAILED* due to *status code*.

Explanation

The servlet received a request containing an undefined action.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

HWNEM0109E Connection test to Element Manager failed.

Explanation

Either the element manager information was entered incorrectly or there was a problem communicating with the element manager.

Action

Check element manager details. Ensure that the element manager is online and that there is a working network connection between both boxes.

HWNEM0111E A problem occurred locating the element manager in IBM Spectrum Control.

Explanation

Element manager not found in the IBM Spectrum Control Database.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

HWNEM0112E A problem occurred updating the element manager to IBM Spectrum Control.

Explanation

Some unexpected error occurred when attempting to update the element manager.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

HWNEM0113E A problem occurred removing the SMI-S provider from the element manager.

Explanation

Some unexpected error occurred when attempting to remove the SMI-S provider.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0114E A problem occurred locating the SMI-S provider associated with the element manager.

Explanation

Some unexpected error occurred when attempting find the element manager's SMI-S provider. It could not be found.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0115E Unable to establish an https connection to the element manager.

Explanation

The servlet received a request containing an undefined action.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0116E A problem occurred removing the element manager from IBM Spectrum Control.

Explanation

Some unexpected error occurred when attempting to remove the element manager.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0117E A problem occurred testing the connection to the element manager.

Explanation

Some unexpected error occurred when attempting to test the connection to the element manager.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0118I Connection test to the element manager *element manager* passed.

Explanation

A connection test was successfully performed on the specified element manager with the displayed result.

Action

The connection test passed and no further action is required.

HWNEM0119E Connection test to the element manager *element manager* failed.

Explanation

A connection could not be established to the specified element manager.

Action

There could be a number of reasons for the failure, including, but not restricted to the element manager's machine being down, network problems or the element manager settings in IBM Spectrum Control.

HWNEM0120E IBM Spectrum Control is unable to communicate with the element manager, the problem could be the element manager is not running, a network communication error or user credentials stored in IBM Spectrum Control are incorrect.

Explanation

A token is required to be exchanged for launching the element manager. No valid token is available. This could be for a number of reasons such as: the element manager might not be running, the network connection might be broken or the element manager user credentials might be wrong.

Action

Verify that the element manager is running at the specified location and that the network connection is good. Check user credentials.

HWNEM0121I Connection test to the SMI-S provider *SMI-S provider* passed.

Explanation

A connection test was successfully performed on the specified SMI-S provider with the published result.

Action

The connection test passed no further action is required.

HWNEM0122I Connection test to the SMI-S provider *SMI-S provider* failed.

Explanation

A connection could not be established to the specified SMI-S provider.

Action

There could be a number of reasons for the failure, including, but not restricted to the SMI-S provider's machine being down, network problems or the SMI-S provider settings in IBM Spectrum Control.

HWNEM0123E Before executing the action the user must select an element manager from the table.

Explanation

The action requires an element manager to execute the action.

Action

Select an element manager from the table.

HWNEM0124E There is no SMI-S provider associated with the selected element manager.

Explanation

The action requires that an element manager has an associated SMI-S provider.

Action

No action necessary since the action is not relevant at this time.

HWNEM0125I The element manager is about to be removed. Once removed the element manager will not be accessible from IBM Spectrum Control. Do you wish to continue.

Explanation

Removes the reference to the element manager from the IBM Spectrum Control. Once the deletion is completed the element manager will not be accessible from IBM Spectrum Control. If access to the element manager is required in future then the element manager needs to be readded.

Action

Click cancel to take no further action and maintain access to the element manager. Click OK to proceed with the removal of the element manager.

HWNEM0126I The element manager is about to be removed. Once removed the element manager will not be accessible from IBM Spectrum Control. Do you wish to continue.

Explanation

Removes the reference to the element manager from the IBM Spectrum Control. Once the deletion is completed the element manager will not be accessible from IBM Spectrum Control. If access to the element manager is required in future, the element manager will need to be re-added.

Action

Click cancel to take no further action and maintain access to the element manager. Click OK to proceed with the removal of the element manager.

HWNEM0127E An internal processing error occurred while servicing the last request.

Explanation

An unexpected error occurred when the backend configuration processed a request.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

HWNEM0128E The element manager is not available. Ensure that the element manager's credentials are defined and up to date.

Explanation

The element manager can not be reached or logged onto.

Action

Ensure that the element manager's user credentials are correct and that the element manager is accessible. If the user credentials are undefined or out of date, update them.

HWNEM0129E An unexpected error occurred changing the element manager's default password.

Explanation

Some unexpected error occurred when the backend configuration processing a request.

Action

This is an internal error. Contact IBM support.

Related reference

-  [Getting support](#)

HWNEM0130E SMI-S provider connection was added to the element manager, however the SMI-S provider discovery job failed to launch.

Explanation

Some unexpected error occurred when the backend configuration processing a request.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0131I The DS8000 SMI-S provider has been added successfully. IBM Spectrum Control has started the discovery job for the DS8000 storage subsystem managed by this SMI-S provider. To check the status of the jobs, go to the IBM Spectrum Control perspective and check the following navigation tree nodes: Administrative Services -> Discovery -> CIMOM

Explanation

The SMI-S provider connection was successfully added and the SMI-S provider discovery job was launched.

Action

No action necessary.

HWNEM0132I SMI-S provider connection was removed from element manager.

Explanation

The SMI-S provider connection was successfully removed.

Action

No action necessary.

HWNEM0133E Internal error occurred, the element manager info for launching the element manager could not be located in the IBM Spectrum Control DB.

Explanation

An unexpected error occurred because expected information was missing from the database.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0134E The user session data is no longer available. To continue restart the IBM Spectrum Control GUI.

Explanation

Some unexpected error occurred that resulted in the required user information being lost.

Action

This is an internal error. Contact IBM support.

Related reference

- [Getting support](#)

HWNEM0135E Unable to contact the device server. It appears to be down.

Explanation

IBM Spectrum Control device server is down.

Action

Restart the device server.

HWNEM0136I The DS8000 Element Manager at *IP Address* has been added successfully.

Explanation

Successfully added the element manager.

Action

None.

HWNEM0137I *service.method* IBM Spectrum Control User *IBM Spectrum Control user* launched DS8000 Element Manager *IP Address* under the alias of *DS8000 user name*.

Explanation

Successfully launched the element manager.

Action

None.

HWNEM0138I The DS8000 SMI-S provider has been modified successfully. IBM Spectrum Control has started the discovery job for the DS8000 storage subsystem managed by this SMI-S provider. To check the status of the jobs, go to the IBM Spectrum Control perspective and check the following navigation tree nodes: Administrative Services -> Data Sources -> Discovery -> CIMOM

Explanation

SMI-S provider connection successfully modified and the SMI-S provider discovery job launched.

Action

No action necessary.

HWNEM0139I The DS8000 SMI-S provider has been modified successfully.

Explanation

SMI-S provider connection successfully modified.

Action

No action necessary.

HWNEM0140W The element manager GUI cannot be accessed because no username and password have been specified by the currently logged-in IBM Spectrum Control user. Element manager credentials are managed on a per-IBM Spectrum Control user basis. The element manager GUI will appear after you provide the correct username and password information.

Explanation

SMI-S provider discovery does not provide the element manager username and password. The message will be displayed if the user attempts to launch the element manager GUI after SMI-S provider discovery without manually updating the element manager's user credentials.

Action

Dismiss the popup dialog. The edit element manager dialog appears. Complete the user name and password fields. On successful update of the fields the element manager is automatically launched.

HWNEM0141E The length of the SMI-S provider's text description is too long. The description should be 255 characters or less.

Explanation

The description length is limited to 255 characters.

Action

Re-enter the text description in 255 or less characters

HWNEM0142E An element manager named *manager.name* already exists.

Explanation

The name of an element manager must be unique across the element managers defined in IBM Spectrum Control. The user has attempted to name an element manager with a name that has already been used.

Action

Choose a new name for the element manager.

HWNEM0143E The DS8000 Element Manager Console is only accessible from IBM Spectrum Control. To access it, open IBM Spectrum Control and switch to the DS8000 Element Manager perspective.

Explanation

The user has attempted to access the DMC perspective panels from an external browser session external to IBM Spectrum Control. The panel is only accessible from within IBM Spectrum Control.

Action

HWNFS - File system monitor messages

- [HWNFS0001E The file system monitor tool does not recognize the command: command.](#)
- [HWNFS0002E The file system monitor tool is unable to read its property file.](#)
- [HWNFS0003W The FileSystems property value property_value has an invalid format.](#)
- [HWNFS0004E The file system monitor tool was unable to initialize access to the IBM Spectrum Control database repository.](#)
- [HWNFS0005E The probe of device_name ended with the error_code error code.](#)
- [HWNFS0006E The FileSystems property is missing in the fsmon.properties file.](#)
- [HWNFS0007E Data cannot be retrieved from the database repository.](#)
- [HWNFS0008E Data about IBM SONAS devices or file systems cannot be retrieved from the database repository.](#)
- [HWNFS0009W Capacity data cannot be collected for the file_system file system.](#)
- [HWNFS0010W No matching file system found in the database repository for file_system.](#)

HWNFS0001E The file system monitor tool does not recognize the command: *command*.

Explanation

The specified command is not supported by the file system monitor tool.

Action

The list of supported commands is displayed with the following command: fsmon help.

HWNFS0002E The file system monitor tool is unable to read its property file.

Explanation

The file system monitor tool encountered an error when trying to read the property file fsmon.properties.

Action

The property file fsmon.properties must exist in the configuration directory of the Device server with appropriate access permissions and it must contain the FileSystems property.

HWNFS0003W The FileSystems property value *property_value* has an invalid format.

Explanation

The FileSystems property value includes a file system specification that cannot be parsed. The file system specification is skipped.

Action

Specify a list of file systems using the correct format. The format for the FileSystems property value is the following:
FileSystems=Device1:FileSystems1,Device2,FileSystem1,...

HWNFS0004E The file system monitor tool was unable to initialize access to the IBM Spectrum Control database repository.

Explanation

The file system monitor tool is unable to retrieve the list of IBM SONAS file systems to monitor from the database repository.

Action

Verify that the database connection property file, tsnmdbparms.properties, exists in the configuration directory for the Device server and that the database repository is running and available.

HWNFS0005E The probe of *device_name* ended with the *error_code* error code.

Explanation

The probe schedule for collecting capacity data about file systems from a IBM SONAS device did not complete.

Action

Verify that the IBM SONAS device can be reached from the Device server using an SSH client. If unable to resolve the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNFS0006E The *FileSystems* property is missing in the *fsmon.properties* file.

Explanation

The file system monitor tool requires the *FileSystems* property to determine which file systems to monitor.

Action

Add the *FileSystems* property to the *fsmon.properties* property file located in the configuration directory for the Device server. If you add the *FileSystems* property but do not define a value for it, all the file systems from all IBM SONAS devices are monitored.

HWNFS0007E Data cannot be retrieved from the database repository.

Explanation

The file system monitor was unable to retrieve file system information from the database repository.

Action

Verify that the IBM Spectrum Control database repository is running and available. If unable to resolve the problem, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNFS0008E Data about IBM SONAS devices or file systems cannot be retrieved from the database repository.

Explanation

There is no data available for IBM SONAS devices or file systems in the database repository.

Action

The file system monitor tool requires that a IBM SONAS device be added to IBM Spectrum Control and probed by a data collection schedule. See the Information Center for information about how to add devices for monitoring and run probes to collect data about those devices.

HWNFS0009W Capacity data cannot be collected for the *file_system* file system.

Explanation

The IBM SONAS device did not report any data for the specified file system. This occurs when the file system is deleted from the IBM SONAS device and that device was not probed after the deletion.

Action

Use the IBM SONAS CLI command lsfs to verify that the file system still exists on the IBM SONAS device.

HWNFS0010W No matching file system found in the database repository for *file_system*.

Explanation

The FileSystems property in the fsmon.properties file specifies a file system that is not found in the IBM Spectrum Control database repository.

Action

Update the FileSystems property in the fsmon.properties file to specify file systems that exist in the database repository. The fsmon.properties file is located in the configuration directory for the Device server.

HWNLM - Planner manager messages

- [HWNLM0700E Storage pool recommendations cannot be given because a storage subsystem is not selected as input.](#)
- [HWNLM0701E The selected planner cannot occur because a storage subsystem is not selected as input.](#)
- [HWNLM0702E The selected planner cannot occur because a computer is not selected as input.](#)
- [HWNLM0703E The selected planner\(s\) cannot occur because a computer and/or volume is not selected as input.](#)
- [HWNLM0704E The selected planner cannot occur because a volume is not selected as input.](#)
- [HWNLM0705E Storage pool recommendation can not be done because there is already a storage volume selected as input.](#)
- [HWNLM0706E The selected planner cannot occur because there is already a volume selected as input.](#)
- [HWNLM0707E The selected planner\(s\) cannot occur because a subsystem or computer is not selected as input.](#)
- [HWNLM0708E The total required space \(total space MB\) exceeds the total available space \(total space MB\) considering all planner inputs. This could be due to insufficient space in the selected storage subsystem's controller\(s\) or pool\(s\), or the selected RAID level does not have enough free space or in case of SVC there is insufficient IO group memory configuration.](#)
- [HWNLM0709E The storage cannot be assigned to the specified hosts due to insufficient LUN addresses.](#)
- [HWNLM0710E The specified total size to be allocated \(total space MB\) is invalid. It must be positive and an integral multiple of 100.](#)
- [HWNLM0711E The specified minimum volume size \(minimum size MB\) is invalid. It must be positive and an integral multiple of 100.](#)
- [HWNLM0712E The specified maximum volume size \(maximum size MB\) is invalid. It must be positive and an integral multiple of 100.](#)
- [HWNLM0713E The total capacity requested cannot be obtained. The allowed minimum size is {0} GB and the allowed maximum size is {1} GB for the given subsystem\(s\). The total size is {2} GB.](#)
- [HWNLM0714E Host Host does not have multipath support because it only has one Fibre Channel port.](#)
- [HWNLM0715E There are no common fabrics that have the minimum number of required paths between the selected servers and the managed storage systems.](#)
- [HWNLM0716E There needs to be at least two common fabrics between host Host and storage subsystem Subsystem in order to use the redundant fabric option.](#)
- [HWNLM0718E The number of paths specified was Paths, but the redundant fabric option requires the number of paths to be an even number of paths.](#)
- [HWNLM0719E Host host does not have a Host Bus Adapter \(HBA\) installed. Please select a host with a HBA installed and try again.](#)
- [HWNLM0720E The Plan failed to generate due to storage subsystem Subsystem not having any volumes identified for Planning use. Please select another storage subsystem and try again.](#)
- [HWNLM0721E A supported multipath driver is not installed on host Host.](#)
- [HWNLM0722E There is an insufficient number of Fibre Channel paths between host Host and storage subsystem Subsystem. Requested Paths paths were requested but there are only Possible Paths paths available.](#)
- [HWNLM0723E The number of paths specified was Paths, but the redundant fabric option requires the minimum number of paths for virtual disks to be four paths or more.](#)
- [HWNLM0724E The number of zones in fabric fabric will be number of zones. This is larger than the max zones maximum number of zones specified.](#)
- [HWNLM0725E Storage pool data does not exist for storage subsystem subsystem. This is either due to not running a storage subsystem probe or not having any fixed block formatted storage pools on the subsystem.](#)
- [HWNLM0726E Performance data does not exist for input storage subsystem\(s\) for the given date range. Either run an IBM Spectrum Control Performance Monitor against the given storage subsystem\(s\) or under the Capacity Planner, select the 'Space Only' Workload Profile option.](#)
- [HWNLM0728E An unexpected internal error occurred. Please contact IBM customer technical support.](#)
- [HWNLM0729E The resulting SAN Planner actions include creating a zoneset on McData fabric fabric WWN. Since there is already an active zone set, please select the 'Use active zone set' option under the Zone Planner and try again.](#)
- [HWNLM0730E SDD version 1.6.2.3 installed on HP host host is not supported by IBM Spectrum Control.](#)
- [HWNLM0731E The computer probe for computer host was incomplete. Please attempt another computer probe and try again.](#)
- [HWNLM0732E DM-Multipath installed on Linux host host does not support the selected Multipath mode.](#)
- [HWNLM0734E The total required IO Group mirroring memory \(total memory KB\) for vdisk mirroring creation exceeds total available mirroring memory \(total memory KB\) available.](#)
- [HWNLM0735E Could not find any subsystem for planning. It is possible that the given subsystems are not detectable.](#)
- [HWNLM0736E The primary, secondary or tertiary storage subsystem subsystem that you specified is not registered with the IBM Spectrum Control-R server.](#)
- [HWNLM0737E Storage Subsystem subsystem is not a supported type for replication operations.](#)
- [HWNLM0738E Replication session with a combination of virtualized and non-virtualized storage volumes is not supported.](#)

- [HWNL0739E Replication session type session is not allowed for the supplied subsystem subsystem of typesubsystemtype.](#)
- [HWNL0740E Limit reached : Number of copy pairs for Replication session type session for the supplied subsystem subsystem.](#)
- [HWNL0741E Replication session type session is not supported.](#)
- [HWNL0742E Format of LSS Property File filepath is invalid.](#)
- [HWNL0743E Namespace in specified LSS is not available.](#)
- [HWNL0744E Unable to find suitable placement for replication storage volumes. plannermsg](#)
- [HWNL0745E Secondary SRG secsrc is either empty or does not contain any valid elements for replication related resource provisioning.](#)
- [HWNL0746E Tertiary SRG tersrc is either empty or does not contain any valid elements for replication related resource provisioning.](#)
- [HWNL0747I Please ensure that proper Replication license and device feature codes are enabled, otherwise the plan execution will fail.](#)
- [HWNL0748E Replication Manager is not installed.](#)
- [HWNL0749E Replication Planner Internal Error.](#)
- [HWNL0750E Replication session with the specified name \(sesname \) already exists. Please use a different name.](#)
- [HWNL0751E No storage system resource can satisfy the provisioning requirements.](#)
- [HWNL0752E The subsystem\(s\) provided in the input do not satisfy the given Thin Provisioning criteria.](#)
- [HWNL0753E The subsystem\(s\) provided in the input do not allow provisioning on volumes or no subsystems were selected.](#)
- [HWNL0754E The subsystem\(s\) provided in the input do not allow provisioning on virtual disks or no subsystems were selected.](#)
- [HWNL0755E The selected virtual disk\(s\) cannot be added in a plan for volumes.](#)
- [HWNL0756E The selected volume\(s\) cannot be added in a plan for virtual disks.](#)
- [HWNL0757I The volume size has been slightly adjusted to meet subsystem requirement.](#)
- [HWNL0758I If a FlashCopy source has multiple targets, an IncrementalFlashCopy relationship can be established with one and only one target.](#)
- [HWNL0759I IncrementalFlashCopy is not available with FlashCopy SE.](#)
- [HWNL0760W Input Volume is already in source role for 12 Flash Copy Sessions, ID = volumeID, Volume Name = volumeName.](#)
- [HWNL0761W Input Volume is already in a target role of Flash Copy Session\(s\), ID = volumeID, Volume Name = volumeName.](#)
- [HWNL0762W Input Volume is already in a target role of Continuous Copy Session\(s\), ID = volumeID, Volume Name = volumeName.](#)
- [HWNL0763W Input Volume is already in a target role of Flash Copy Session\(s\).](#)
- [HWNL0764I Ensure connectivity between source and target \(direct or through fabric\).](#)
- [HWNL0765W For TSE volumes, please ensure that the Repository Capacity is configured and available on the pool poolName.](#)
- [HWNL0766W The storage subsystem ssName is unacceptable due to -- reason.](#)
- [HWNL0767I The storage subsystem ssName is a valid candidate subsystem, thus it is considered during planning.](#)
- [HWNL0768W The storage pool poolName is unacceptable due to -- reason.](#)
- [HWNL0769I The storage pool poolName is a valid candidate storage pool, thus it is considered during planning.](#)
- [HWNL0770W The SVC ssName is unacceptable due to -- reason.](#)
- [HWNL0771I The SVC ssName is a valid candidate subsystem, thus it is considered during planning.](#)
- [HWNL0772W The mdiskgroup mdiskGroupName is unacceptable due to -- reason.](#)
- [HWNL0773I The mdiskgroup mdiskGroupName is a valid candidate mdiskgroup, thus it is considered during planning.](#)
- [HWNL0774W The iogroup ioGroupName is unacceptable due to -- reason.](#)
- [HWNL0775I The iogroup ioGroupName is a valid candidate iogroup, thus it is considered during planning.](#)
- [HWNL0776E Unable to plan for volumeName due to reaching the maximum limit of volumes on the LSS\(es\) within storage pool storagePool on storage subsystem subsystem.](#)
- [HWNL0777E Replication session with the specified name \(sesname \) associated with SRG \(srgname \) has required replication and does not require extension.](#)
- [HWNL0778I The volume/vdisk name may be different during plan execution based on the name availability and/or subsystem limitation.](#)
- [HWNL0779E Unable to recommend plan due to reaching the maximum limit of volumes on the LSS\(es\) within storage pool\(s\) storagePool on storage subsystem\(s\) subsystem.](#)
- [HWNL0780E Unable to recommend plan since the LSS range specified in the LSSRange.properties file is not valid for volumeName on storage subsystem subsystem.](#)
- [HWNL0781E No cluster partnership exists between source subsystem srcSS and target subsystem targetSS.](#)
- [HWNL0782E No connectivity path exists between source subsystem srcSS and target subsystem targetSS.](#)
- [HWNL0783E The selected input volume volumeName is missing from the storage subsystem.](#)
- [HWNL0784E The selected input virtual disk virtualDiskName is missing from the SVC.](#)
- [HWNL0785E Provisioning with replication can not be achieved with Extent Space Efficient\(ESE\) or Track Space Efficient\(TSE\) volumes. Extent Space Efficient volumes are not allowed in copy sets and Track Space Efficient volumes are valid only in the target role of a FlashCopy session or the Journal role of a Global Mirror or Metro Global Mirror session.](#)
- [HWNL0786E Replication can not be extended to Extent Space Efficient\(ESE\) or Track Space Efficient\(TSE\) source volumes. Extent Space Efficient volumes are not allowed in copy sets and Track Space Efficient volumes are valid only in the target role of a FlashCopy session or the Journal role of a Global Mirror or Metro Global Mirror session. One or more selected input volumes are either ESE or TSE : volumes.](#)
- [HWNL0787E vDisk Mirroring is already enabled for the input vdisk volumes.](#)
- [HWNL0788I Fabric selection is based on planner selection of the host port\(s\) and the subsystem port\(s\) and not by the fabric\(s\) selected by the user.](#)
- [HWNL0789E The IO group IO Group for virtual disk virtual disk does not have appropriate connectivity to the host\(s\) selected in the plan.](#)
- [HWNL0790E Not considering Subsystem\(s\): Subsystem because not all of its back-end subsystems are configured into this IBM Spectrum Control or some are currently undetectable. This is unacceptable when an input RAID level is specified or a workload profile other than space-only is selected.](#)
- [HWNL0791E The input volume you specified is already defined as a snapshot copy volume of a snapshot copy session, volume ID = volume_id, volume name = volume_name.](#)
- [HWNL0792E The input volume you specified cannot be defined as a source volume in a replication session, ID = volume_id, volume name = volume_name.](#)
- [HWNL0793I Replication Manager is used to manage snapshot copy sessions.](#)
- [HWNL0794E The total space requirement \(total space MB\) cannot be met within a single storage pool for the replication session. This could be due to insufficient space in the selected storage subsystem's pool\(s\), or because the storage subsystem's pool\(s\) do not meet the requirements of the Planner input.](#)
- [HWNL0796E The input volumes you specified must be within a single storage pool for the replication session.](#)
- [HWNL0797E Volumes could not be created or used in the source pool of the existing snapshot copy session.](#)
- [HWNL0798E The total space requirement \(total space MB\) cannot be met within the storage pool of the existing replication session. This could be due to insufficient space in the storage pool of the existing replication session, or because the input volumes you specified are not within the storage pool of the existing replication session.](#)
- [HWNL0799E The input volumes you specified are already defined in an existing snapshot copy session.](#)
- [HWNL0800W The input volume you specified is already defined as a snapshot copy volume of a snapshot copy session, volume ID = volume_id, volume name = volume_name.](#)
- [HWNL0801E The input volumes you specified are not within the storage pool of the existing replication session.](#)
- [HWNL0802E For a Metro Global Mirror session, you must specify at least three DS8000 subsystems with sufficient capacity.](#)
- [HWNL0803E The selected input volume volume_name is already in a copy set used by this copy session.](#)

- [HWNLM0804E Subsystem cannot be considered. Please check both the candidate SVC and its back-end subsystems for available space and performance data for the specified time interval.](#)
- [HWNLM0804I The storage pool\(s\) provided do not satisfy the criteria for an acceptable destination pool.](#)
- [HWNLM0805E Host Host is a virtual machine without any Fibre Channel host port and a storage volume cannot be assigned directly to it.](#)
- [HWNLM0806W Host Host is a hypervisor. IBM Spectrum Control is unable to set the multipath policy on the hypervisor.](#)
- [HWNLM0807W Host Host is a server that is not managed by an SRA. IBM Spectrum Control is unable to set the multipath policy on the host.](#)
- [HWNLM0808E Ports of host Host are not connected to any fabric that is known to IBM Spectrum Control.](#)
- [HWNLM0809E Ports of host Host are not connected to fabrics that allow automatic zoning.](#)
- [HWNLM0810E The total capacity requested is {0} GB for {1} number of volumes. The individual volume size comes out to be {2} GB, which is invalid. The allowed minimum size is {3} GB and the allowed maximum size is {4} GB per volume for the given subsystem\(s\).](#)
- [HWNLM0811W No fabric information is available. All fabric-related options are ignored and no fabric configuration operation is performed.](#)
- [HWNLM0812E Unable to plan virtualizer provisioning task due to reaching the maximum vdisks limitation in all candidate iogroups.](#)
- [HWNLM0001I An integrated SAN Planner job started with schedule creator. schedule name](#)
- [HWNLM0002E The integrated SAN Planner job completed with errors. Message from exception: message.](#)
- [HWNLM0003I The integrated SAN Planner job completed.](#)
- [HWNLM0004W The integrated SAN Planner job completed with warnings.](#)
- [HWNLM0005E The integrated SAN Planner job completed with errors.](#)
- [HWNLM0006I Zone set zone set name created on fabric fabric wwn.](#)
- [HWNLM0007I Zone zone name created on fabric fabric wwn.](#)
- [HWNLM0008I Zone zone name added to zone set zone set name on fabric fabric wwn.](#)
- [HWNLM0009I A list of ports added to zone zone name for zone set zone set name on fabric fabric wwn.](#)
- [HWNLM0010I Activated zone set zone set name on fabric zone set name.](#)
- [HWNLM0011I Started to create storage volumes.](#)
- [HWNLM0012E The creation of storage volumes completed with errors.](#)
- [HWNLM0013I Completed creating storage volumes.](#)
- [HWNLM0014I Started to assign storage volumes to WWPNs.](#)
- [HWNLM0015E The assignment of storage volumes to WWPNs completed with errors.](#)
- [HWNLM0016I Completed assigning storage volumes to WWPNs.](#)
- [HWNLM0017W The command to discover volumes on host host id failed with status status.](#)
- [HWNLM0018W Unable to set the multipath policy on host host id due to host failure status.](#)
- [HWNLM0019I Completed startTransaction command on fabric fabric wwn.](#)
- [HWNLM0020I Completed commitTransaction command on fabric fabric wwn.](#)
- [HWNLM0021E The startTransaction command on fabric fabric wwn failed with return code return code.](#)
- [HWNLM0022E Creation for Zone set zone set name on fabric fabric wwn failed with return code return code.](#)
- [HWNLM0023E Creation for Zone zone name created on fabric fabric wwn failed with return code return code.](#)
- [HWNLM0024E Adding Zone zone name to zone set zone set name on fabric fabric wwn failed with return code return code.](#)
- [HWNLM0025E Adding ports to zone zone name for zone set zone set name on fabric fabric wwn failed with return code return code.](#)
- [HWNLM0026E Activated zone set zone set name on fabric zone set name failed with return code return code.](#)
- [HWNLM0027E The commitTransaction command on fabric fabric wwn failed with return code return code.](#)
- [HWNLM0028I Starting volume discovery on host host.](#)
- [HWNLM0029I Finished volume discovery on host host.](#)
- [HWNLM0030I Assignment\(s\) between Volume volume id and Host Port\(s\) host ports already exist, no assignment actions will happen for these paths.](#)
- [HWNLM0031W Since multiple Storage Resource Groups were provided as input to the plan, the newly created volumes will not be added to any Storage Resource Groups.](#)
- [HWNLM0032W IBM Spectrum Control is unable to set the multipath policy on host host id because it is an ESX hypervisor. After the provisioning operation completes, log in to the hypervisor or use VMware tools to set the multipath policy on the hypervisor.](#)
- [HWNLM0033W IBM Spectrum Control is unable to set the multipath policy on host host id because it is not managed by an SRA. After the volume or volumes are provisioning, log in to the operating system on the agentless server to set the multipath policy.](#)
- [HWNLM0034I Started updating agentless server configuration with disk information.](#)
- [HWNLM0035I Completed updating agentless server information with disk information.](#)
- [HWNLM0036W The volume discovery operation failed for one or more agentless servers.](#)
- [HWNLM0037I Disks mappings detected for volume volume by the discovery on host host on the following paths: path.](#)
- [HWNLM0038W No disks were detected for volume volume by the discovery on host host.](#)
- [HWNLM0039I After unassigning volume volume from host host, no disks were detected anymore for it by the discovery on that host.](#)
- [HWNLM0040W After unassigning volume volume from host host, disks mappings were still detected for it by the discovery on that host on the following paths: path.](#)
- [HWNLM0041I Started to copy storage volumes.](#)
- [HWNLM0042E The copy of storage volumes could not be completed.](#)
- [HWNLM0043I Completed copy of storage volumes.](#)
- [HWNLM0100E No Storage Subsystem\(s\) passed to SAN Planner.](#)
- [HWNLM0101E No Storage Virtualized Controller\(s\) passed to SAN Planner.](#)
- [HWNLM0102E Virtual disk\(s\) were selected in a plan for volumes.](#)
- [HWNLM0103E Volume\(s\) were selected in a plan for virtual disks.](#)
- [HWNLM0104I Not considering subsystem {0} for new storage because the user indicated so.](#)
- [HWNLM0105I Not considering storage pool {0} for new storage because the user indicated so.](#)
- [HWNLM0106I Not considering storage virtualized controller {0} for new storage because the user indicated so.](#)
- [HWNLM0107I Not considering managed disk group {0} for new storage because the user indicated so.](#)
- [HWNLM0108I Considering virtual disk: {0} of storage virtualized controller: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship \(or these requirements were over-ridden by user\).](#)
- [HWNLM0109I Considering user selected virtual disk: {0} of storage virtualized controller: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship \(or these requirements were over-ridden by user\).](#)
- [HWNLM0110I Considering user selected volume: {0} of subsystem: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship \(or these requirements were over-ridden by user\).](#)
- [HWNLM0111I The algorithm used to identify the best location for the volumes will ignore the co-location criterion.](#)
- [HWNLM0300I A Path Planner job started with multipath policy \(multipath\), redundant fabric policy \(rFabric\), multipath mode \(mode\), and \(paths\) number of paths.](#)
- [HWNLM0301I The Path Planner job completed.](#)
- [HWNLM0302E Multiple paths are not supported on host host name as it has only one port.](#)
- [HWNLM0303E No common fabrics between Host and Subsystem](#)

- [HWNLM0304E There is an insufficient number of possible paths between Host and Subsystem. The number of possible paths possiblePaths are less than the required number of paths paths.](#)
- [HWNLM0305E Cannot create redundant paths using the specified paths number of paths.](#)
- [HWNLM0306E There are less than two fabrics in common between host Host and storage subsystem Subsystem.](#)
- [HWNLM0307W No supported multipath driver was found on host host id.](#)
- [HWNLM0308W No corresponding multipath device was found on host host id in case of one or more volumes.](#)
- [HWNLM0309W Multipath policy configuration is not supported for multipath driver on host host id.](#)
- [HWNLM0310W Failed to set the multipath policy on host host id because for Multipath DM driver only Round Robin policy is available.](#)
- [HWNLM0311W The command for setting the multipath policy on host host id has failed to execute.](#)
- [HWNLM0312I Please check the agent log file for more details.](#)
- [HWNLM0313W iSCSI ports were not mapped to the storage volume because storage subsystem Subsystem does not support iSCSI connectivity.](#)
- [HWNLM0314W iSCSI ports were not mapped to the storage volume because storage subsystem Subsystem does not have iSCSI connectivity configured.](#)
- [HWNLM0315E No common connectivity exists between Host and Subsystem.](#)
- [HWNLM0316I Started to set the multipath policy on host host id.](#)
- [HWNLM0317I The multipath policy has been successfully set on host host id.](#)
- [HWNLM0500I The user specified maximum number of zones is user zones however the current number of zones is current zones.](#)
- [HWNLM0501E The maximum number of zones in a fabric policy was violated.](#)
- [HWNLM0502E The zone per host policy was violated.](#)
- [HWNLM0503E The zone per HBA policy was violated.](#)
- [HWNLM0504E The zone per cluster policy was violated.](#)
- [HWNLM0505E The zone per controller policy was violated.](#)
- [HWNLM0506E The zone per controller type policy was violated.](#)
- [HWNLM0507E The zone per fabric policy was violated.](#)
- [HWNLM0508I Zone Set zone set name was created.](#)
- [HWNLM0509I Zone zone name was created.](#)
- [HWNLM0510I Zone zone name was added to Zone Set zone set name.](#)
- [HWNLM0511I Host port port id was added to Zone zone name.](#)
- [HWNLM0512I Subsystem port port id was added to Zone zone name.](#)
- [HWNLM0513I SAN Planner started with guidance Policy \(guidance\), validation policies \(validation\), fabric WWN \(fabricWWN\), and using active zone set \(zone set name\).](#)
- [HWNLM0514I The Zone Planner completed.](#)
- [HWNLM0515W Not considering Subsystem {0} since Planner could not obtain its information from database \(it could be undetectable\).](#)
- [HWNLM0516E Could not find any subsystem with the given input. The subsystem\(s\) may be undetectable.](#)
- [HWNLM0517W Invalid candidate subsystem {0} for performance data for SAN Planner.](#)
- [HWNLM0518I Not considering Pool: {0} of Subsystem: {1} since it does not have enough allocatable space.](#)
- [HWNLM0519W Not considering Pool: {0} of Subsystem: {1} since it is not online.](#)
- [HWNLM0520I Not considering Pool: {0} of Subsystem: {1} because it is not Fixed Block pool.](#)
- [HWNLM0521I Not considering Pool: {0} of Subsystem: {1} because it is Solid State Disk pool.](#)
- [HWNLM0522I Not considering Pool: {0} of Subsystem: {1} because it is not a Solid State Disk pool.](#)
- [HWNLM0523I Not considering Pool: {0} of Subsystem: {1} because its Encryption Group is not matching with the input.](#)
- [HWNLM0524I Not considering Pool: {0} of Subsystem: {1} because it is not thin-provisioning enabled.](#)
- [HWNLM0525I Not considering Pool: {0} of Subsystem: {1} because its lock behavior is not matching with the input.](#)
- [HWNLM0526I Not considering Pool: {0} of Subsystem: {1} because it is not in the selected input pools list or it is already filtered out.](#)
- [HWNLM0527I Not considering Pool: {0} of Subsystem: {1} because its RAID level does not match with the input RAID level.](#)
- [HWNLM0528I Considering volume: {0} of Subsystem: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship \(or these requirements were over-ridden by user\).](#)
- [HWNLM0529W The subsystem {0} does not have performance data for the specified time interval.](#)
- [HWNLM0530I Not considering volume {0} for new storage because the user indicated so.](#)
- [HWNLM0531I Not considering the subsystem {0} because it does not support Extent Space Efficient volumes.](#)
- [HWNLM0532I Not considering the Subsystem: {0} because it does not support Track Space Efficient volumes.](#)
- [HWNLM0533I Not considering Pool: {0} of Subsystem: {1} because it does not have repository capacity available/defined for Track Space Efficient volumes.](#)
- [HWNLM0534I New Capacity Planning Advice Task Started....](#)
- [HWNLM0535I Capacity Planning Advice Task Completed.](#)
- [HWNLM0536I Not considering Pool: {0} of Subsystem: {1} because its backend storage RAID level does not match the input RAID level](#)
- [HWNLM0537I Not considering Pool: {0} of Subsystem: {1} because its backend storage is not configured into this IBM Spectrum Control or is currently undetectable and input requires a specific backend RAID level](#)
- [HWNLM0538I Not considering IO Group: {0} of Subsystem: {1} because it does not have two nodes associated with it](#)
- [HWNLM0539I Not considering IO Group: {0} of Subsystem: {1} because it has more than or equal to 2048 vdisks](#)
- [HWNLM0540I Not considering IO Group: {0} of Subsystem: {1} because it does not have enough available mirroring memory](#)
- [HWNLM0541I Not considering Volume: {0} of Subsystem: {1} because its subsystem was found to be unacceptable based on inputs](#)
- [HWNLM0542I Not considering Volume: {0} of Subsystem: {1} because its pool was found to be unacceptable based on inputs](#)
- [HWNLM0543I Not considering Volume: {0} of Subsystem: {1} because its IO group was found to be unacceptable based on inputs](#)
- [HWNLM0544I Not considering Volume: {0} of Subsystem: {1} because its size does not match the input](#)
- [HWNLM0545I Not considering Subsystem: {0} because not all of its backend subsystems are configured into this IBM Spectrum Control or some are currently undetectable. This is unacceptable when an input RAID level is specified or a workload profile other than space-only is selected](#)
- [HWNLM0546I Not considering Subsystem: {0} because IBM Spectrum Control does not have adequate performance data for it and/or some of its backend subsystems. For planning with workload profiles other than space-only, this subsystem and all of its backend subsystems need to have daily performance data](#)
- [HWNLM0547I Not considering Subsystem: {0} since none of its IO groups have enough available mirroring memory to support vdisk mirroring input](#)
- [HWNLM0548I Not considering Subsystem: {0} since it does not support vdisk mirroring \(code level is below v4.3\)](#)
- [HWNLM0549I Not considering Subsystem: {0} since it does not support space-efficient vdisks \(code level is below v4.3\)](#)
- [HWNLM0550I Not considering IO Group: {0} of Subsystem: {1} because it does not appropriate connectivity to hosts selected in the input](#)
- [HWNLM0551I Not considering Subsystem: {0} since none of its IO groups have appropriate connectivity to hosts selected in the input](#)
- [HWNLM0552I Not considering Volume: {0} of Subsystem: {1} since it does not match input on thin-provisioning characteristics](#)
- [HWNLM0553I Volume size too small {0} MB. Minimum size should be {1} MB for the pool {2}. Not considering this pool.](#)
- [HWNLM0554I Volume size too big {0} MB. Maximum size should be {1} MB for the pool {2}. Not considering this pool.](#)
- [HWNLM0555I Not considering Volume: {0} of Subsystem: {1} because it is in a known replication relationship.](#)
- [HWNLM0556I Volume {0} can not be moved to the following candidate pool\(s\) due to insufficient allocatable space: {1}.](#)
- [HWNLM0557I The max theoretical I/O capability of storage pool {0} on subsystem {1} has not been set. Using default value {2}.](#)
- [HWNLM0558W The subsystem pool {0} does not have performance data for the specified time interval.](#)

- [HWNLM0559I Not considering IO Group: {0} of Subsystem: {1} because compression was specified, but the IO Group does not have compression active.](#)
- [HWNLM0560I Not considering Volume: {0} of Subsystem {1} since it is not compressed when compression was specified.](#)
- [HWNLM0560W Fabric agent not available to perform the zoning operation on the fabric {0}.](#)
- [HWNLM0561I Fabric agent available to perform the zoning operation on the fabric {0}.](#)
- [HWNLM0562E Fabric service exception occurred trying to check for zone control on the fabric {0}.](#)
- [HWNLM0563W A zone unnecessary zone was not created, because the host already has connectivity to the storage system using existing zone.](#)
- [HWNLM0564W There is already a zone named zone_name. A zone named new_zone_name was created instead.](#)
- [HWNLM0565W There is already a volume named vol_name. A volume named new_vol_name will be created instead.](#)
- [HWNLM0566I Not considering Pool: {0} of Subsystem: {1} because it is mixed pool.](#)
- [HWNLM0567I Not considering Subsystem: {0} for vdisk mirroring since it does not have at least two acceptable mdiskgroups from different backend subsystems.](#)
- [HWNLM0568I Not considering I/O group {0} of subsystem {1} because it already reached the limit of 200 compressed vdisks.](#)
- [HWNLM0569I Not considering subsystem {0} because none of its I/O groups can be used for compressed vdisks.](#)
- [HWNLM0570I Not considering I/O group {0} of subsystem {1} because host definition {2} has access restrictions.](#)
- [HWNLM0571I Not considering storage system {0} because none of its I/O groups allow restricted access to all hosts.](#)
- [HWNLM801I Synchronous Refresh Configuration of the Storage Subsystem was completed Successfully for Subsystem subsystem.](#)
- [HWNLM802E Synchronous Refresh Configuration of the Storage Subsystem did not complete successfully for Subsystem subsystem.](#)
- [HWNLM803I Replication Session was created successfully session.](#)
- [HWNLM804E Replication Session creation failed. ReplicationManager \[session\].](#)
- [HWNLM805I Successful check for existence of session session.](#)
- [HWNLM806E No Such session exists with name session.](#)
- [HWNLM807I CopySets Added to Session successfully session.](#)
- [HWNLM808E CopySets addition to session failed. ReplicationManager \[session\].](#)
- [HWNLM809E CopySets creation failed. ReplicationManager \[session\].](#)
- [HWNLM810I Storage Subsystem Configuration refreshed successfully in Replication Manager ss.](#)
- [HWNLM811E Storage Subsystem Configuration refresh operation failed in Replication Manager \[ss\].](#)
- [HWNLM812E Fabric agent not available to perform the zoning operation on the fabric {0}.](#)
- [HWNLM813I Replication Session was started successfully session.](#)
- [HWNLM814W Replication Session start failed. ReplicationManager \[session\].](#)
- [HWNLM815W Not considering the Storage Subsystem {0} because it is not registered in TPC-R.](#)
- [HWNLM816I Please use Replication Manager console to schedule point-in-time copy creation.](#)
- [HWNLM817W No cluster partnership exists between source subsystem \[srcss\] and target subsystem \[tgtss\].](#)
- [HWNLM818W No path exists between source subsystem \[srcss\] and target subsystem \[tgtss\].](#)
- [HWNLM820I Copysets can not be added to a session when it's in Prepared/Suspend/Recover/Flash state. Please try again later after some time. State of session can be checked from Replication Manager console.](#)
- [HWNLM0821W IBM Spectrum Control had no information about the Fibre Channel configuration, so it did not verify fabric connectivity or change the zoning configuration.](#)

HWNLM0700E Storage pool recommendations cannot be given because a storage subsystem is not selected as input.

Explanation

In order to recommend storage pools, a storage subsystem must first be selected as input.

Action

Select a storage subsystem and try again.

HWNLM0701E The selected planner cannot occur because a storage subsystem is not selected as input.

Explanation

In order to continue, a storage subsystem must first be selected as input.

Action

Select a storage subsystem and try again.

HWNLM0702E The selected planner cannot occur because a computer is not selected as input.

Explanation

In order to continue, a computer must first be selected as input.

Action

Select a computer and try again.

HWNLM0703E The selected planner(s) cannot occur because a computer and/or volume is not selected as input.

Explanation

In order to continue, a computer and/or storage volume must first be selected as input.

Action

Select a computer and/or storage volume and try again.

HWNLM0704E The selected planner cannot occur because a volume is not selected as input.

Explanation

In order to continue, a storage volume must first be selected as input.

Action

Select a storage volume and try again.

HWNLM0705E Storage pool recommendation can not be done because there is already a storage volume selected as input.

Explanation

In order to recommend storage pools, a storage volume must not be selected as input.

Action

Unselect the storage volume as input and try again.

HWNLM0706E The selected planner cannot occur because there is already a volume selected as input.

Explanation

In order to continue, a storage volume must not be selected as input.

Action

Unselect the storage volume as input and try again.

HWNLM0707E The selected planner(s) cannot occur because a subsystem or computer is not selected as input.

Explanation

In order to continue, a computer and/or storage subsystem must first be selected as input.

Action

Select a computer and/or storage subsystem and try again.

HWNLMO708E The total required space (*total space* MB) exceeds the total available space (*total space* MB) considering all planner inputs. This could be due to insufficient space in the selected storage subsystem's controller(s) or pool(s), or the selected RAID level does not have enough free space or in case of SVC there is insufficient IO group memory configuration.

Explanation

Based on the Planner input, there is not enough free space for the request.

Action

Either reduce the required space or select another storage subsystem or pool with enough free space. It is also possible that there were not enough volumes available in the LSS(es) of the ESS subsystem(s) selected. If this is the case, then ensure that there are enough volumes available in the ESS subsystem LSS(es) to accommodate the request.

HWNLMO709E The storage cannot be assigned to the specified hosts due to insufficient LUN addresses.

Explanation

The number of LUN addresses on the host is at its maximum.

Action

Select another host to assign the storage to and try again.

HWNLMO710E The specified total size to be allocated (*total space* MB) is invalid. It must be positive and an integral multiple of 100.

Explanation

The total size specified is an invalid number.

Action

Correct the total size value and try again.

HWNLMO711E The specified minimum volume size (*minimum size* MB) is invalid. It must be positive and an integral multiple of 100.

Explanation

The minimum volume size specified is an invalid number.

Action

Correct the minimum size value and try again.

HWNLMO712E The specified maximum volume size (*maximum size* MB) is invalid. It must be positive and an integral multiple of 100.

Explanation

The maximum volume size specified is an invalid number.

Action

Correct the maximum size value and try again.

HWNLMO713E The total capacity requested cannot be obtained. The allowed minimum size is {0} GB and the allowed maximum size is {1} GB for the given subsystem(s). The total size is {2} GB.

Explanation

The specified total size is invalid for the given storage subsystem(s).

Action

Correct the invalid volume size for the given subsystem(s).

HWNLMO714E Host *Host* does not have multipath support because it only has one Fibre Channel port.

Explanation

In order for the host to have multipath support, there needs to be at least two Fibre Channel ports on the host.

Action

Select a host capable of multipath support and try again.

HWNLMO715E There are no common fabrics that have the minimum number of required paths between the selected servers and the managed storage systems.

Explanation

In order for the SAN Planner action to proceed, there needs to be at least one common fabric between the hosts and storage subsystems. If multipath planning is specified, at least two data paths between each host and storage subsystem is required. For virtual disks, an I/O group requires two nodes to be connected to a common fabric. If only one node of the I/O group has a working connection to the fabric, then the I/O group cannot be used for path planning.

Action

Select one or more hosts and storage subsystems within the same fabric, then try the SAN Planner action again. For virtual disks, check the status of the connectivity of the nodes of the I/O groups on the SVC. Ensure both nodes of each I/O group have a working connection to the same fabrics as the selected hosts, then probe the SVC and fabrics again. Once the probes complete, try the SAN Planner action again.

HWNLMO716E There needs to be at least two common fabrics between host *Host* and storage subsystem *Subsystem* in order to use the redundant fabric option.

Explanation

In order to use the redundant fabric option, there needs to be at least two common fabrics between the host and storage subsystem. For virtual disks, an I/O group requires two nodes to be connected to at least two common fabrics. If only one node of the I/O group has a working connection to the fabric, then the I/O group cannot be used for path planning.

Action

Either clear the redundant fabric option or select a host and storage subsystem that have two fabrics common between them and try again. For virtual disks, check the status of the connectivity of the nodes of the I/O groups on the SVC. Ensure both nodes of each I/O group have a working connection to at least two common fabric(s) as the selected host(s) and then reprobe the SVC and the fabric(s). Once the probes complete, try again.

HWNLM0718E The number of paths specified was *Paths*, but the redundant fabric option requires the number of paths to be an even number of paths.

Explanation

The redundant fabric option requires the number of paths to be an even number of paths.

Action

Alter the number of paths so that the number of paths is even or clear the redundant fabric option. Once the alteration has been completed, the plan can be submitted again.

HWNLM0719E Host *host* does not have a Host Bus Adapter (HBA) installed. Please select a host with a HBA installed and try again.

Explanation

The Path Planner requires Fibre Channel connectivity between the host and the storage subsystem. Without a Host Bus Adapter installed this connectivity cannot be made and storage subsystem volumes cannot be assigned to the host.

Action

Select another host or install a Host Bus Adapter on this host and try again.

HWNLM0720E The Plan failed to generate due to storage subsystem *Subsystem* not having any volumes identified for Planning use. Please select another storage subsystem and try again.

Explanation

Host volume assignment requires at least one volume to be identified for use by the SAN Planner.

Action

Modify the planner input and try again.

HWNLM0721E A supported multipath driver is not installed on host *Host*.

Explanation

The selected host does not have a supported multipath driver installed. Currently, the only supported multipath driver is IBM SDD.

Action

Select a host with a supported multipath driver and try again.

HWNLM0722E There is an insufficient number of Fibre Channel paths between host *Host* and storage subsystem *Subsystem*. Requested *Paths*

paths were requested but there are only *Possible Paths* paths available.

Explanation

The number of Fibre Channel paths available is less than the user requested number of paths.

Action

Either add more Fibre Channel paths or decrease the number of required paths and try again.

HWNLM0723E The number of paths specified was *Paths*, but the redundant fabric option requires the minimum number of paths for virtual disks to be four paths or more.

Explanation

The redundant fabric option requires the minimum number of paths for virtual disks to be four paths or more.

Action

Alter the number of paths for virtual disks so that the number of paths is even and equal or greater than four paths or uncheck the redundant fabric option. Once the alteration has been completed, the plan can be submitted again.

HWNLM0724E The number of zones in fabric *fabric* will be *number of zones*. This is larger than the *max zones* maximum number of zones specified.

Explanation

If this Planner action would continue, the resulting number of zones would exceed the user specified maximum number of zones.

Action

Either increase the maximum number of zones or change the Planner inputs and try again.

HWNLM0725E Storage pool data does not exist for storage subsystem *subsystem*. This is either due to not running a storage subsystem probe or not having any fixed block formatted storage pools on the subsystem.

Explanation

In order for the Volume Planner to continue, it requires the selected storage subsystem to have a valid storage pool.

Action

First verify the storage subsystem has been probed. If it has been probed, verify the storage subsystem's storage pools are fixed block formatted and not CKD formatted.

HWNLM0726E Performance data does not exist for input storage subsystem(s) for the given date range. Either run an IBM Spectrum Control Performance Monitor against the given storage subsystem(s) or under the Capacity Planner, select the 'Space Only' Workload Profile option.

Explanation

The Capacity Planner requires performance data in order to determine what the best plan of action will be for the selected storage subsystem and Planner options. Without performance data, you can still proceed by selecting the 'Space Only' Workload Profile option. With this option selected, the Capacity Planner will make all decisions based solely on the storage capacity levels. For SVC storage subsystem, it is required that the performance data is collected for all the back-end storage subsystems of the SVC as well as the performance data collected for the SVC itself.

Action

Either run an IBM Spectrum Control Performance Monitor against this storage subsystem or select the 'Space Only' Workload Profile option under the Capacity Planner and try again.

HWNLM0728E An unexpected internal error occurred. Please contact IBM customer technical support.

Explanation

The SAN Planner job failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNLM0729E The resulting SAN Planner actions include creating a zoneset on McData fabric *fabric WWN*. Since there is already an active zone set, please select the 'Use active zone set' option under the Zone Planner and try again.

Explanation

Since IBM Spectrum Control 3.3 supports only inband fabric agents for McData, there is a current limitation that only one zoneset can be defined for the fabric. An existing zone set must be deleted before a new one can be created, otherwise newly created zone set will replace the current active zone set.

Action

Select the 'Use active zone set' option under the zone planner and try again, so SAN Planner will add the recommended zone(s) into current active zone set.

HWNLM0730E SDD version 1.6.2.3 installed on HP host *host* is not supported by IBM Spectrum Control.

Explanation

HP SDD version 1.6.2.3 has a problem which will cause the SAN Planner action to fail.

Action

A problem with the SDD 'cfgvpath -r' command causes a failure. This is specific to SDD version 1.6.2.3 on HP-UX only. Please upgrade to a different SDD version and try again.

HWNLM0731E The computer probe for computer *host* was incomplete. Please attempt another computer probe and try again.

Explanation

The SAN Planners were unable to obtain multipath driver information obtained from a computer probe.

Action

An unknown problem caused the computer probe to be incomplete. Check the computer probe's job log to determine the problem and try again.

HWNLM0732E DM-Multipath installed on Linux host *host* does not support the selected Multipath mode.

Explanation

The selected host has DM-Multipath installed which only supports the round-robin mode.

Action

DM-Multipath only supports the round-robin mode. Please select this multipath option.

HWNLM0734E The total required IO Group mirroring memory (*total memory* KB) for vdisk mirroring creation exceeds total available mirroring memory (*total memory* KB) available.

Explanation

Based on the Planner input, there is not enough free IO Group mirroring memory for the request.

Action

Either reduce the required space or select another IBM SAN Volume Controller with enough free IO Group mirroring memory.

HWNLM0735E Could not find any subsystem for planning. It is possible that the given subsystems are not detectable.

Explanation

Based on the Planner input, the subsystems provide are not detectable in IBM Spectrum Control, possibly in missing state, hence planner can not use them.

Action

Run the probe again and check the status of the subsystem in IBM Spectrum Control.

HWNLM0736E The primary, secondary or tertiary storage subsystem *subsystem* that you specified is not registered with the IBM Spectrum Control-R server.

Explanation

The storage subsystem that you specified as a primary, secondary or tertiary candidate is not registered with the IBM Spectrum Control for Replication server. Depending on the replication session type that you specified, there must be at least one primary, one secondary, and and one tertiary storage subsystem registered with the IBM Spectrum Control for Replication server to generate a recommendation.

Action

Depending on the replication session type that you specified, ensure that the primary, secondary and tertiary storage subsystem candidates exist and then register the primary, secondary and tertiary storage subsystem with the IBM Spectrum Control for Replication server. After the storage subsystems have been added, refresh the IBM Spectrum Control for Replication configuration.

HWNLM0737E Storage Subsystem *subsystem* is not a supported type for replication operations.

Explanation

Storage Subsystem(s) are not supported for replication operations.

Action

Check the types of supported subsystems for replication operation.

HWNL0738E Replication session with a combination of virtualized and non-virtualized storage volumes in not supported.

Explanation

Input contains a mix of virtualized and non-virtualized storage volumes. Replication session with a combination of virtualized and non-virtualized storage volumes in not supported.

Action

Avoid mixing virtualized and non-virtualized storage volumes in one replication session.

HWNL0739E Replication session type *session* is not allowed for the supplied subsystem *subsystem* of type *subsystemtype*.

Explanation

Planner input contains type of subsystems that does not support specified replication session type.

Action

Planner input contains type of subsystems that does not support specified replication session type.

HWNL0740E Limit reached : Number of copy pairs for Replication session type *session* for the supplied subsystem *subsystem*.

Explanation

Limit reached for number of copy pairs of particular session type of supplied storage subsystem.

Action

Limit reached for number of copy pairs of particular session type of supplied storage subsystem.

HWNL0741E Replication session type *session* is not supported.

Explanation

Input replication session type is not supported.

Action

Input replication session type is not supported.

HWNL0742E Format of LSS Property File *filepath* is invalid.

Explanation

Format of LSS Property File is invalid.

Action

Format of LSS Property File is invalid.

HWNLM0743E Namespace in specified LSS is not available.

Explanation

Namespace in specified LSS is not available. Check the LSS Property File.

Action

Namespace in specified LSS is not available. Check the LSS Property File.

HWNLM0744E Unable to find suitable placement for replication storage volumes. *plannermsg*

Explanation

Unable to find suitable placement for replication storage volumes. Readjust input requirements and(or) candidate entities. This may happen because of various reasons like: There was no subsystem provided for secondary/tertiary volumes, or The target subsystem is not eligible for planning based on the input provided in the planner (for example Thin Provisioning, Solid State Drives etc), or No eligible pools were found on the target subsystem, or The subsystem is not added to IBM Spectrum Control for Replication, or In case of San Volume Controller, there is no connectivity established between the source and the target subsystems, or The input contains a number of SAN Volume Controller subsystems and some of which have connectivity established, but some of them do not have the connectivity established between themselves, etc.

Action

Unable to find suitable placement for replication storage volumes. Readjust input requirements and(or) candidate entities.

HWNLM0745E Secondary SRG *secsrg* is either empty or does not contain any valid elements for replication related resource provisioning.

Explanation

Secondary SRG is either empty or does not contain any valid elements for replication related resource provisioning.

Action

Check secondary SRG contents to validate the requirements.

HWNLM0746E Tertiary SRG *tersrg* is either empty or does not contain any valid elements for replication related resource provisioning.

Explanation

Tertiary SRG is either empty or does not contain any valid elements for replication related resource provisioning.

Action

Check tertiary SRG contents to validate the requirements.

HWNLM0747I Please ensure that proper Replication license and device feature codes are enabled, otherwise the plan execution will fail.

Explanation

The SAN Planner does not check for license on IBM Spectrum Control for Replication and the feature codes on the storage subsystems. Please ensure that proper Replication license and device feature codes are enabled, otherwise the plan execution will fail.

Action

The SAN Planner does not check for license on IBM Spectrum Control for Replication and the feature codes on the storage subsystems. Please ensure that proper Replication license and device feature codes are enabled, otherwise the plan execution will fail.

HWNLMO748E Replication Manager is not installed.

Explanation

Replication Manager is not installed.

Action

Install Replication Manager.

HWNLMO749E Replication Planner Internal Error.

Explanation

Replication Planner Internal Error. Unable to retrieve information from Replication Manager.

Action

Replication Planner Internal Error. Unable to retrieve information from Replication Manager.

HWNLMO750E Replication session with the specified name (*sesname*) already exists. Please use a different name.

Explanation

Replication session with the specified name already exists. Please use a different name.

Action

Replication session with the specified name already exists. Please use a different name.

HWNLMO751E No storage system resource can satisfy the provisioning requirements.

Explanation

No storage system resource can satisfy the requested storage capacity and the requirements of the service class.

Action

Try again, specifying different requirements. If you are provisioning from a capacity pool, consider adding storage systems to the capacity pool, selecting a different capacity pool, or allowing provisioning from any capacity pool or any available storage. Consider selecting a different service class.

HWNLMO752E The subsystem(s) provided in the input do not satisfy the given Thin Provisioning criteria.

Explanation

The given subsystem(s) do not support Thin Provisioning or the given Thin Provisioning method.

Action

Try again with different subsystem(s) or with changed input parameters. The SAN Planner log may provide more information.

HWNLM0753E The subsystem(s) provided in the input do not allow provisioning on volumes or no subsystems were selected.

Explanation

The selection to provision volumes will not work against selected the subsystem(s) because the selected subsystem(s) do not allow the provisioning of volumes or there were no subsystem(s) selected in the plan.

Action

Either change the provisioning request to virtual disks or add subsystem(s) that allow volume provisioning.

HWNLM0754E The subsystem(s) provided in the input do not allow provisioning on virtual disks or no subsystems were selected.

Explanation

The selection to provision virtual disks will not work against selected the subsystem(s) because the selected subsystem(s) do not allow the provisioning of virtual disks or there were no subsystem(s) selected in the plan.

Action

Either change the provisioning request to volumes or add subsystem(s) that allow virtual disk provisioning.

HWNLM0755E The selected virtual disk(s) cannot be added in a plan for volumes.

Explanation

The selection to plan for volumes will not work against the selected virtual disk(s).

Action

Either remove the selected virtual disk(s) or remove the selected volume(s) and submit the plan for virtual disks.

HWNLM0756E The selected volume(s) cannot be added in a plan for virtual disks.

Explanation

The selection to plan for virtual disks will not work against the selected volume(s).

Action

Either remove the selected volume(s) or remove the selected virtual disk(s) and submit the plan for volumes.

HWNLM0757I The volume size has been slightly adjusted to meet subsystem requirement.

Explanation

The subsystem allows creation of volumes in multiples of a specific number, which varies for different subsystems. The individual volume size is therefore adjusted by the SAN Planner for the given subsystem(s).

Action

HWNLM0758I If a FlashCopy source has multiple targets, an IncrementalFlashCopy relationship can be established with one and only one target.

Explanation

An incremental flashCopy relationship can be established with one and only one target.

Action

HWNLM0759I IncrementalFlashCopy is not available with FlashCopy SE.

Explanation

Action

HWNLM0760W Input Volume is already in source role for 12 Flash Copy Sessions, ID = *volumeID*, Volume Name = *volumeName*.

Explanation

Action

HWNLM0761W Input Volume is already in a target role of Flash Copy Session(s), ID = *volumeID*, Volume Name = *volumeName*.

Explanation

Action

HWNLM0762W Input Volume is already in a target role of Continuous Copy Session(s), ID = *volumeID*, Volume Name = *volumeName*.

Explanation

Action

HWNLM0763W Input Volume is already in a target role of Flash Copy Session(s).

Explanation

Action

HWNLMO764I Ensure connectivity between source and target (direct or through fabric).

Explanation

Action

HWNLMO765W For TSE volumes, please ensure that the Repository Capacity is configured and available on the pool *poolName*.

Explanation

Action

HWNLMO766W The storage subsystem *ssName* is unacceptable due to -- *reason*.

Explanation

Action

HWNLMO767I The storage subsystem *ssName* is a valid candidate subsystem, thus it is considered during planning.

Explanation

Action

HWNLMO768W The storage pool *poolName* is unacceptable due to -- *reason*.

Explanation

Action

HWNLMO769I The storage pool *poolName* is a valid candidate storage pool, thus it is considered during planning.

Explanation

Action

HWNLMO770W The SVC *ssName* is unacceptable due to -- *reason*.

Explanation

Action

HWNLM0771I The SVC *ssName* is a valid candidate subsystem, thus it is considered during planning.

Explanation

Action

HWNLM0772W The mdiskgroup *mdiskGroupName* is unacceptable due to -
- *reason*.

Explanation

Action

HWNLM0773I The mdiskgroup *mdiskGroupName* is a valid candidate mdiskgroup, thus it is considered during planning.

Explanation

Action

HWNLM0774W The iogroup *ioGroupName* is unacceptable due to --
reason.

Explanation

Action

HWNLM0775I The iogroup *ioGroupName* is a valid candidate iogroup, thus it is considered during planning.

Explanation

Action

HWNLM0776E Unable to plan for *volumeName* due to reaching the maximum limit of volumes on the LSS(es) within storage pool *storagePool* on storage subsystem *subsystem*.

Explanation

The maximum number of volumes on each LSS within a storage pool is 256. The plan recommendation found that of the eligible storage pools, there are no longer LSS(es) available to create a storage volume on the storage subsystem since creating a volume on such a storage pool will result in exceeding the maximum number of volumes on

the LSS(es) within the storage pool.

Action

Either change the provisioning request to select storage pool(s) where the LSS(es) have not reached the maximum number of volumes or select storage subsystem(s) where the maximum number volumes on the LSS(es) within the storage pools has yet to be reached. Also note that the LSS range for the volume role may be changed in the LSSRange.properties file located in the device/conf directory where the product was installed. Increasing the range may help depending on the type of DS8000 storage volume being created.

HWNLM0777E Replication session with the specified name (*sesname*) associated with SRG (*srgname*) has required replication and does not require extension.

Explanation

SRG and associated session has required replication. It does not require extension.

Action

SRG and associated session has required replication. It does not require extension.

HWNLM0778I The volume/vdisk name may be different during plan execution based on the name availability and/or subsystem limitation.

Explanation

SAN Planner gives the volume name recommendation, but it is possible that at the time of execution of the plan, the volume name was already taken or the subsystem does not allow user specific names, so the resulting volume name will be different.

Action

HWNLM0779E Unable to recommend plan due to reaching the maximum limit of volumes on the LSS(es) within storage pool(s) *storagePool* on storage subsystem(s) *subsystem*.

Explanation

The maximum number of volumes on each LSS within a storage pool is 256. The plan recommendation found that of the eligible storage pools, there are no longer LSS(es) available to create a storage volume on the storage subsystem since creating a volume on such a storage pool will result in exceeding the maximum number of volumes on the LSS(es) within the storage pool.

Action

Either change the provisioning request to select storage pool(s) where the LSS(es) have not reached the maximum number of volumes or select storage subsystem(s) where the maximum number volumes on the LSS(es) within the storage pools has yet to be reached.

HWNLM0780E Unable to recommend plan since the LSS range specified in the LSSRange.properties file is not valid for *volumeName* on storage subsystem *subsystem*.

Explanation

The LSSRange.properties file is required to process storage volumes on DS8000. This file is located under the device/conf directory where the product is installed and identifies the LSS range for each type of volume that is to be created by the plan. The plan encountered an error accessing this file.

Action

Ensure the LSSRange.properties file exists and the entries inside the file are in the proper format for the DS8000 volumes to be created. The range of eligible LSS(es) is in the format of 0Xnn-0Xmm where nn and mm represent the hexadecimal digit of the LSS lower and upper boundaries. The lower bound nn must be less than or equal to the upper bound mm. A full explanation of valid entries exist within the LSSRange.properties file.

HWNLMO781E No cluster partnership exists between source subsystem *srcSS* and target subsystem *targetSS*.

Explanation

No cluster partnership exists between source subsystem and target subsystem.

Action

HWNLMO782E No connectivity path exists between source subsystem *srcSS* and target subsystem *targetSS*.

Explanation

No connectivity path exists between source subsystem and target subsystem.

Action

HWNLMO783E The selected input volume *volumeName* is missing from the storage subsystem.

Explanation

The selected input volume is missing from the storage subsystem and is required to complete the requested plan. The volume may have been deleted from the storage subsystem or the volume may have not been discovered after a scheduled storage subsystem probe.

Action

Remove the selected input volume from the plan or choose to ignore the selected input volume so that it is not considered in the plan recommendation. If the selected input volume exists in a storage resource group, remove the volume from the storage resource group or choose to ignore the selected input volume in the storage resource group.

HWNLMO784E The selected input virtual disk *virtualDiskName* is missing from the SVC.

Explanation

The selected input virtual disk is missing from the SVC and is required to complete the requested plan. The virtual disk may have been deleted from the SVC or the virtual disk may have not been discovered after a scheduled SVC probe.

Action

Remove the selected input virtual disk from the plan or choose to ignore the selected input virtual disk so that it is not considered in the plan recommendation. If the selected input virtual disk exists in a storage resource group, remove the virtual disk from the storage resource group or choose to ignore the selected input virtual disk in the storage resource group.

HWNLMO785E Provisioning with replication can not be achieved with Extent Space Efficient(ESE) or Track Space Efficient(TSE) volumes. Extent Space Efficient volumes are not allowed in copy sets and Track Space Efficient volumes are valid only in the target role of

a FlashCopy session or the Journal role of a Global Mirror or Metro Global Mirror session.

Explanation

Extent Space Efficient volumes are not allowed in copy sets, Track Space Efficient volumes are valid only in the target role of a FlashCopy session or the Journal role of a Global Mirror or Metro Global Mirror session.

Action

Please use standard volumes or use a thin provisioning profile which meets this limitation.

HWNLMO786E Replication can not be extended to Extent Space Efficient(ESE) or Track Space Efficient(TSE) source volumes. Extent Space Efficient volumes are not allowed in copy sets and Track Space Efficient volumes are valid only in the target role of a FlashCopy session or the Journal role of a Global Mirror or Metro Global Mirror session. One or more selected input volumes are either ESE or TSE : *volumes*.

Explanation

Extent Space Efficient volumes are not allowed in copy sets, Track Space Efficient volumes are valid only in the target role of a FlashCopy session or the Journal role of a Global Mirror or Metro Global Mirror session.

Action

Please use standard volumes or use thin provisioned volumes that meets this limitation.

HWNLMO787E vDisk Mirroring is already enabled for the input vdisk *volumes*.

Explanation

vDisk mirroring is already enabled for the input vDisk.

Action

vDisk mirroring is already enabled for the input vDisk.

HWNLMO788I Fabric selection is based on planner selection of the host port(s) and the subsystem port(s) and not by the fabric(s) selected by the user.

Explanation

SAN Planner will ignore the fabrics selected by the user and select the fabric(s) based on the port(s) that are selected when SAN planner finds the best single or multiple paths from the host port(s) to the subsystem port(s). The zone changes for the fabric(s) are then based on the fabric(s) of the ports that have been selected by the SAN Planner.

Action

HWNLMO789E The IO group *IO Group* for virtual disk *virtual disk* does not have appropriate connectivity to the host(s) selected in the plan.

Explanation

The IO group that connects the virtual disk to the fabric does not have enough fabrics or enough paths to satisfy the request to assign the virtual disks to the host(s) specified in the plan.

Action

If redundant fabrics have not been selected, reduce the number of paths to at least two paths or ensure the IO group nodes have connectivity to the same fabric as the host(s). If redundant fabrics have been selected, reduce the number of paths to at least four paths or ensure the specified IO group has two nodes connected to different fabrics shared by the host(s).

HWNLMO790E Not considering Subsystem(s): *Subsystem* because not all of its back-end subsystems are configured into this IBM Spectrum Control or some are currently undetectable. This is unacceptable when an input RAID level is specified or a workload profile other than space-only is selected.

Explanation

Not considering the current subsystem(s) because not all of its back-end subsystems are configured into this IBM Spectrum Control or some are currently undetectable. This is unacceptable when an input RAID level is specified or a workload profile other than space-only is selected.

Action

HWNLMO791E The input volume you specified is already defined as a snapshot copy volume of a snapshot copy session, volume ID = *volume_id*, volume name = *volume_name*.

Explanation

A volume cannot be the source volume of a replication session if it is already defined as a snapshot copy volume.

Action

Select a different volume to use as the source volume.

HWNLMO792E The input volume you specified cannot be defined as a source volume in a replication session, ID = *volume_id*, volume name = *volume_name*.

Explanation

A volume cannot be the source volume of a replication session if it is already defined as a snapshot copy volume or if it is not defined to the Replication Manager.

Action

Select a different volume to use as the source volume.

HWNLMO793I Replication Manager is used to manage snapshot copy sessions.

Explanation

The Planner recommends the source volumes for the snapshot copy session, but it does not control the creation of snapshots. Management of snapshot copies is performed using Replication Manager.

Action

Use Replication Manager to manage snapshot copy sessions.

HWNLMO794E The total space requirement (*total space* MB) cannot be met within a single storage pool for the replication session. This could be due to insufficient space in the selected storage subsystem's pool(s), or because the storage subsystem's pool(s) do not meet the requirements of the Planner input.

Explanation

Based on the Planner input, there is not enough space in a single storage pool for the replication session request.

Action

Either reduce the required space or select another storage subsystem or pool with enough free space to meet the requirements of the Planner input.

HWNLMO796E The input volumes you specified must be within a single storage pool for the replication session.

Explanation

The input volumes for a replication session must reside within the same storage pool to generate a valid replication session.

Action

Create a plan for a replication session that will only specify volumes within the same storage pool.

HWNLMO797E Volumes could not be created or used in the source pool of the existing snapshot copy session.

Explanation

The required volumes could not be created or used in the source pool of the existing snapshot copy session.

Action

Create a plan for an existing snapshot copy session that will satisfy the request in the source pool of the existing snapshot copy session.

HWNLMO798E The total space requirement (*total space* MB) cannot be met within the storage pool of the existing replication session. This could be due to insufficient space in the storage pool of the existing replication session, or because the input volumes you specified are not within the storage pool of the existing replication session.

Explanation

Based on the Planner input, there is not enough space in the storage pool of the existing replication session or the input volumes specified are not within the storage pool of the existing replication session.

Action

Create a plan for a replication session that will only specify volumes with the same storage pool as the existing replication session or reduce the required space.

HWNLM0799E The input volumes you specified are already defined in an existing snapshot copy session.

Explanation

The input volumes for a snapshot copy session must not already be defined in an existing snapshot copy session.

Action

Create a plan for a snapshot copy session that will specify volumes not already defined in an existing snapshot copy session.

HWNLM0800W The input volume you specified is already defined as a snapshot copy volume of a snapshot copy session, volume ID = *volume_id*, volume name = *volume_name*.

Explanation

The input volume for a continuous copy session are already defined as a snapshot copy volume of a snapshot copy session.

Action

Select a different volume to use as the source volume.

HWNLM0801E The input volumes you specified are not within the storage pool of the existing replication session.

Explanation

The input volumes specified are not within the storage pool of the existing replication session.

Action

Create a plan for a replication session that will only specify volumes with the same storage pool as the existing replication session.

HWNLM0802E For a Metro Global Mirror session, you must specify at least three DS8000 subsystems with sufficient capacity.

Explanation

Based on the Planner input, there are not enough DS8000 subsystems with sufficient capacity to create all source and target volumes required for Metro Global Mirror.

Action

Select at least three DS8000 subsystems with sufficient capacity.

HWNLM0803E The selected input volume *volume_name* is already in a copy set used by this copy session.

Explanation

An input volume is already in use by a copy set that is part of this copy session. A volume may not be added to the same session more than once.

Action

Remove this volume from either the planner input or from the existing copy session.

HWNLMO804E Subsystem cannot be considered. Please check both the candidate SVC and its back-end subsystems for available space and performance data for the specified time interval.

Explanation

Either candidate SVC or its back-end subsystems (or both) do not have either space or performance data for the specified time interval.

Action

Make sure that both candidate SVC and its back-end subsystems (or both) have space and performance data for the specified time interval.

HWNLMO804I The storage pool(s) provided do not satisfy the criteria for an acceptable destination pool.

Explanation

The reason the storage pool(s) is not considered acceptable is logged in the log file.

Action

Try again with different storage pool(s) or with changed input parameters. The log file may provide more information.

HWNLMO805E Host *Host* is a virtual machine without any Fibre Channel host port and a storage volume cannot be assigned directly to it.

Explanation

A storage volume cannot be assigned directly to a virtual machine if the virtual machine has no Fibre Channel host port. Instead, it must be assigned to the hypervisor that contains the virtual machine.

Action

Specify the hypervisor for the provisioning operation. If the provisioning operation succeeds, use VMware tools to assign the hypervisor disk to the virtual machine.

HWNLMO806W Host *Host* is a hypervisor. IBM Spectrum Control is unable to set the multipath policy on the hypervisor.

Explanation

IBM Spectrum Control can create volumes, complete fabric zoning, and assign volumes for the hypervisor. However, IBM Spectrum Control cannot enable the multipath policy for the hypervisor.

Action

After the provisioning operation completes, log in to the hypervisor or use VMware tools to enable the multipath policy on the hypervisor.

HWNLMO807W Host *Host* is a server that is not managed by an SRA. IBM Spectrum Control is unable to set the multipath policy on the host.

Explanation

IBM Spectrum Control can create volumes, complete fabric zoning, and assign volumes for the agentless server. However, IBM Spectrum Control cannot enable the multipath policy for the agentless server.

Action

After the provisioning operation completes, log in to the operating system on the agentless server to enable the multipath policy.

HWNLM0808E Ports of host *Host* are not connected to any fabric that is known to IBM Spectrum Control.

Explanation

The host that you are attempting to provision volumes to does not appear to IBM Spectrum Control to have fabric connectivity. It is likely that the host ports are connected to a fabric, but that the fabric has not been probed.

Action

Probe the fabrics to which the host is connected, and repeat the provisioning operation.

HWNLM0809E Ports of host *Host* are not connected to fabrics that allow automatic zoning.

Explanation

Automatic zoning is enabled, but the fabrics to which the host is connected do not allow the automatic zoning operations. For example, this error is returned if the host is connected only to a Brocade fabric that is controlled only by an inband agent.

Action

If the host already has fabric connectivity to the storage system through existing zones, disable automatic zoning and repeat the provisioning operation.

HWNLM0810E The total capacity requested is {0} GB for {1} number of volumes. The individual volume size comes out to be {2} GB, which is invalid. The allowed minimum size is {3} GB and the allowed maximum size is {4} GB per volume for the given subsystem(s).

Explanation

The specified total volume size is invalid for the given number of volumes.

Action

Correct the invalid volume size or the number of volumes and try again.

HWNLM0811W No fabric information is available. All fabric-related options are ignored and no fabric configuration operation is performed.

Explanation

Typically, you cannot provision volumes to servers or hypervisors that do not have fabric connectivity. However, because all of the selected hosts have Fibre Channel Port WWPNs and none appear to IBM Spectrum Control to have fabric connectivity, it is possible that the fabrics were not probed. For this reason, you can provision volumes even though the hosts appear to have no fabric connectivity. However, because no fabric information is available, all of the fabric-related options (number of paths, redundant fabrics, and automatic zoning) are ignored and no fabric configuration operation is performed.

Action

Ensure that the hosts are connected to the recommended storage system before you run the task. After the provisioning operation completes, you must manually configure the fabric.

HWNLM0812E Unable to plan virtualizer provisioning task due to reaching the maximum vdisks limitation in all candidate iogroups.

Explanation

All iogroups on the predefined for the provisioning task virtualizer storage have exceeded the allowed number of vdisk per iogroup.

Action

Please ensure that the predefined for the provisioning task virtualizer storage has any iogroups with less than maximum allowed volumes

HWNLM0001I An integrated SAN Planner job started with schedule *creator. schedule name*

Explanation

None.

Action

None.

HWNLM0002E The integrated SAN Planner job completed with errors. Message from exception: *message*.

Explanation

The SAN Planner job failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

-  [Getting support](#)

HWNLM0003I The integrated SAN Planner job completed.

Explanation

None.

Action

None.

HWNLM0004W The integrated SAN Planner job completed with warnings.

Explanation

None.

Action

None.

HWNLM0005E The integrated SAN Planner job completed with errors.

Explanation

The SAN Planner job failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

-  [Getting support](#)

HWNLM0006I Zone set *zone set name* created on fabric *fabric wwn*.

Explanation

None.

Action

None.

HWNLM0007I Zone *zone name* created on fabric *fabric wwn*.

Explanation

None.

Action

None.

HWNLM0008I Zone *zone name* added to zone set *zone set name* on fabric *fabric wwn*.

Explanation

None.

Action

None.

HWNLM0009I A list of ports added to zone *zone name* for zone set *zone set name* on fabric *fabric wwn*.

Explanation

None.

Action

None.

HWNL0010I Activated zone set zone set name on fabric zone set name .

Explanation

None.

Action

None.

HWNL0011I Started to create storage volumes .

Explanation

None.

Action

None.

HWNL0012E The creation of storage volumes completed with errors .

Explanation

The creation of storage volumes failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

-  [Getting support](#)
-

HWNL0013I Completed creating storage volumes .

Explanation

None.

Action

None.

HWNL0014I Started to assign storage volumes to WWPNS .

Explanation

None.

Action

None.

HWNL0015E The assignment of storage volumes to WWPNS completed with errors .

Explanation

The assignment of storage volumes to WWPNs failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

HWNL00016I Completed assigning storage volumes to WWPNs.

Explanation

None.

Action

None.

HWNL00017W The command to discover volumes on host *host id* failed with status *status*.

Explanation

The discovery of volumes on the host failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

HWNL00018W Unable to set the multipath policy on host *host id* due to host failure *status*.

Explanation

Setting the multipath policy on the host failed due to a variety of reasons. Detail of the failure is logged by the host. The host must be installed with a multipath device driver fully supported by IBM Spectrum Control, but may fail due to the following reasons:

- A multipath driver fully supported by IBM Spectrum Control is not installed on the host.
- A multipath driver fully supported by IBM Spectrum Control does not support the requested multipath policy.
- A multipath driver fully supported by IBM Spectrum Control does not support the assigned subsystem volumes.
- The host failed to fully complete the request on a multipath driver fully supported by IBM Spectrum Control.

Action

Check that you are using a supported multipath driver. For information on driver support, search the IBM Knowledge Center for "multipath subsystem device drivers".

If the multipath driver fully supported by IBM Spectrum Control is installed on the host and meets the conditions, then it will be necessary to investigate messages issued by the host for more details. If the multipath driver fully supported by the IBM Tivoli Storage Productivity Center is not installed on the host, then the multipath policy will need to be set on the host with a multipath driver that supports changing the multipath policy.

HWNL00019I Completed startTransaction command on fabric *fabric wwn*.

Explanation

None.

Action

None.

HWNL0020I Completed commitTransaction command on fabric *fabric wwn*.

Explanation

None.

Action

None.

HWNL0021E The startTransaction command on fabric *fabric wwn* failed with return code *return code*.

Explanation

The command failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNL0022E Creation for Zone set *zone set name* on fabric *fabric wwn* failed with return code *return code*.

Explanation

The zone set creation failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNL0023E Creation for Zone *zone name* created on fabric *fabric wwn* failed with return code *return code*.

Explanation

The zone creation failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNL0024E Adding Zone *zone name* to zone set *zone set name* on fabric *fabric wwn* failed with return code *return code*.

Explanation

Adding zone to the zone set failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNL0025E Adding ports to zone *zone name* for zone set *zone set name* on fabric *fabric wwn* failed with return code *return code*.

Explanation

Adding port to the zone failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNL0026E Activated zone set *zone set name* on fabric *zone set name* failed with return code *return code*.

Explanation

Activating the zone set failed due to an unexpected error.

Action

Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNL0027E The `commitTransaction` command on fabric *fabric wwn* failed with return code *return code*.

Explanation

Provisioning was not successful because of an unexpected error. This problem might occur when a provisioning job is run at the same time as other zone control sessions are in progress for a fabric.

Action

Ensure that no other zone control sessions for the fabric are in progress when the Provisioning job is run. Other zone control sessions might be run from another installation of IBM Spectrum Control, from the command-line interface for a switch, or from the element manager for a switch.

If the fabric is a Brocade fabric and you have a licensed version of Network Advisor, check the Network Advisor event logs for a history of other zone control sessions.

If the fabric is a Cisco fabric, you might need to distribute the full zone database from one of the logical switches to the other logical switches in the VSAN.

If this problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNL0028I Starting volume discovery on host *host*.

Explanation

None.

Action

None.

HWNL0029I Finished volume discovery on host *host*.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNL0030I Assignment(s) between Volume *volume id* and Host Port(s) *host ports* already exist, no assignment actions will happen for these paths.

Explanation

None.

Action

None.

HWNL0031W Since multiple Storage Resource Groups were provided as input to the plan, the newly created volumes will not be added to any Storage Resource Groups.

Explanation

Newly created volumes will automatically be added to a Storage Resource Group only if a single Storage Resource Group is provided as input to the plan.

Action

If you want newly created volumes to automatically be placed into the Storage Resource Group, then use only one Storage Resource Group as input to the SAN Planner.

HWNL0032W IBM Spectrum Control is unable to set the multipath policy on host *host id* because it is an ESX hypervisor. After the provisioning operation completes, log in to the hypervisor or use VMware tools to set the multipath policy on the hypervisor.

Explanation

IBM Spectrum Control can create volumes, complete fabric zoning, and assign volumes for the hypervisor. However, IBM Spectrum Control cannot enable the multipath policy for the hypervisor.

Action

After the provisioning operation completes, log in to the hypervisor or use VMware tools to enable the multipath policy on the hypervisor.

HWNLM0033W IBM Spectrum Control is unable to set the multipath policy on host *host id* because it is not managed by an SRA. After the volume or volumes are provisioning, log in to the operating system on the agentless server to set the multipath policy.

Explanation

IBM Spectrum Control can create volumes, complete fabric zoning, and assign volumes for the agentless server. However, IBM Spectrum Control cannot enable the multipath policy for the agentless server.

Action

After the provisioning operation completes, log in to the operating system on the agentless server to enable the multipath policy.

HWNLM0034I Started updating agentless server configuration with disk information.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNLM0035I Completed updating agentless server information with disk information.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNLM0036W The volume discovery operation failed for one or more agentless servers.

Explanation

IBM Spectrum Control was unable to update the database with information about the volume or volumes that were provisioned to agentless servers.

Action

Run storage system probes for the storage systems on which the volumes are located. Running the probes updates the data for the agentless servers that have volumes that are assigned from the storage systems.

HWNL00037I Disks mappings detected for volume *volume* by the discovery on host *host* on the following paths: *path*.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNL00038W No disks were detected for volume *volume* by the discovery on host *host*.

Explanation

The volume was assigned to the host, but SRA discovered no disks corresponding to this volume on the host.

Action

A disk might be detected by the host later, so a probe that runs later might discover the disk that is related to the volume.

Check for other causes that would prevent the host from detecting the disk. For example, make sure that the driver is configured correctly and that the host has connectivity with the storage system.

HWNL00039I After unassigning volume *volume* from host *host*, no disks were detected anymore for it by the discovery on that host.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNL00040W After unassigning volume *volume* from host *host*, disks mappings were still detected for it by the discovery on that host on the following paths: *path*.

Explanation

The volume was unassigned from the host, but SRA still discovered disks corresponding to this volume on the host.

Action

You can try to run the SRA probe for the host again.

HWNL00041I Started to copy storage volumes.

Explanation

No further information is available.

Action

No further action is required.

HWNLMO042E The copy of storage volumes could not be completed.

Explanation

The copy of storage volumes failed with errors.

Action

Check the log files for the external process and contact IBM Software Support.

HWNLMO043I Completed copy of storage volumes.

Explanation

No further information is available.

Action

No further action is required.

HWNLMO100E No Storage Subsystem(s) passed to SAN Planner.

Explanation

None.

Action

None.

HWNLMO101E No Storage Virtualized Controller(s) passed to SAN Planner.

Explanation

None.

Action

None.

HWNLMO102E Virtual disk(s) were selected in a plan for volumes.

Explanation

None.

Action

None.

HWNLMO103E Volume(s) were selected in a plan for virtual disks.

Explanation

None.

Action

None.

HWNLMO104I Not considering subsystem {0} for new storage because the user indicated so.

Explanation

None.

Action

None.

HWNLMO105I Not considering storage pool {0} for new storage because the user indicated so.

Explanation

None.

Action

None.

HWNLMO106I Not considering storage virtualized controller {0} for new storage because the user indicated so.

Explanation

None.

Action

None.

HWNLMO107I Not considering managed disk group {0} for new storage because the user indicated so.

Explanation

None.

Action

None.

HWNLMO108I Considering virtual disk: {0} of storage virtualized controller: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship (or these requirements were over-ridden by user).

Explanation

None.

Action

None.

HWNLMO109I Considering user selected virtual disk: {0} of storage virtualized controller: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship (or these requirements were over-ridden by user).

Explanation

None.

Action

None.

HWNLMO110I Considering user selected volume: {0} of subsystem: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship (or these requirements were over-ridden by user).

Explanation

None.

Action

None.

HWNLMO111I The algorithm used to identify the best location for the volumes will ignore the co-location criterion.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNLMO300I A Path Planner job started with multipath policy (*multipath*), redundant fabric policy (*rFabric*), multipath mode (*mode*), and (*paths*) number of paths.

Explanation

None.

Action

None.

HWNLMO301I The Path Planner job completed.

Explanation

None.

Action

None.

HWNL0302E Multiple paths are not supported on host *host name* as it has only one port.

Explanation

The host requires two or more HBA ports in order for multiple paths to be support

Action

Modify the planner input and try again.

HWNL0303E No common fabrics between *Host* and *Subsystem*

Explanation

There are no common fabrics connecting the host and subsystem.

Action

Modify the planner input and try again.

HWNL0304E There is an insufficient number of possible paths between *Host* and *Subsystem*. The number of possible paths *possiblePaths* are less than the required number of paths *paths*.

Explanation

The total number of possible paths is less than the required number of paths.

Action

Modify the planner input and try again.

HWNL0305E Cannot create redundant paths using the specified *paths* number of paths.

Explanation

Redundant fabrics requires an even number of paths.

Action

Modify the planner input and try again.

HWNL0306E There are less than two fabrics in common between host *Host* and storage subsystem *Subsystem*.

Explanation

Redundant fabrics requires 2 or more fabrics linking each host with each subsystem.

Action

Modify the planner input and try again.

HWNLMO307W No supported multipath driver was found on host *host id*.

Explanation

No supported multipath driver is installed on this system.

Action

Make sure that fully supported multipath driver is available on this system for Path Planner configurations and try again.

HWNLMO308W No corresponding multipath device was found on host *host id* in case of one or more volumes.

Explanation

For one or more volumes, no matching SDD device has been identified on the host.

Action

Make sure that the volumes have been assigned and they are visible to the host.

HWNLMO309W Multipath policy configuration is not supported for multipath driver on host *host id*.

Explanation

Multipath policy configuration is not supported for multipath driver this host.

Action

Make sure that the multipath driver supports the requested multipath policy.

HWNLMO310W Failed to set the multipath policy on host *host id* because for Multipath DM driver only Round Robin policy is available.

Explanation

For Multipath DM driver only Round Robin policy is available.

Action

If different multipath policy is required, then a supported multipath driver that supports this policy is required.

HWNLMO311W The command for setting the multipath policy on host *host id* has failed to execute.

Explanation

The command for configuring the multipath policy on the host has failed to execute.

Action

Refer to the agent log file for more details.

HWNL0312I Please check the agent log file for more details.

Explanation

The agent log file located on the agent host may contain more details about the problem.

Action

Refer to the agent log file for more details.

HWNL0313W iSCSI ports were not mapped to the storage volume because storage subsystem *Subsystem* does not support iSCSI connectivity.

Explanation

iSCSI ports were not mapped to the storage volume because the storage subsystem does not support iSCSI connectivity..

Action

Access the storage volume using Fibre Channel connectivity. If iSCSI access is required, the volume should be moved to a storage subsystem supporting iSCSI connectivity.

HWNL0314W iSCSI ports were not mapped to the storage volume because storage subsystem *Subsystem* does not have iSCSI connectivity configured.

Explanation

The storage subsystem does not have any target iSCSI ports configured.

Action

Configure the storage subsystem to add iSCSI target ports.

HWNL0315E No common connectivity exists between *Host* and *Subsystem*.

Explanation

The host is configured only for iSCSI connectivity and the subsystem either does not support iSCSI connectivity or does not have iSCSI connectivity configured.

Action

If the subsystem supports iSCSI connectivity configure the storage subsystem for iSCSI connectivity and try again. If the subsystem does not support iSCSI connectivity, either add HBA cards to the host for Fibre Channel access or move the volume to a storage system which is capable of iSCSI connectivity.

HWNL0316I Started to set the multipath policy on host *host id*.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNLM0317I The multipath policy has been successfully set on host *host id*.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNLM0500I The user specified maximum number of zones is *user zones* however the current number of zones is *current zones*.

Explanation

None.

Action

None.

HWNLM0501E The maximum number of zones in a fabric policy was violated.

Explanation

The number of zones the Zone Planner generated are more than the maximum number of zones that the user specified.

Action

Modify the planner input and try again.

HWNLM0502E The zone per host policy was violated.

Explanation

The Zone Planner cannot satisfy the zone per host guidance policy.

Action

Modify the planner input and try again.

HWNLM0503E The zone per HBA policy was violated.

Explanation

The Zone Planner cannot satisfy the zone per HBA guidance policy.

Action

Modify the planner input and try again.

HWNLM0504E The zone per cluster policy was violated.

Explanation

The Zone Planner cannot satisfy the zone per cluster guidance policy.

Action

Modify the planner input and try again.

HWNLM0505E The zone per controller policy was violated.

Explanation

The Zone Planner cannot satisfy the zone per controller guidance policy.

Action

Modify the planner input and try again.

HWNLM0506E The zone per controller type policy was violated.

Explanation

The Zone Planner cannot satisfy the zone per controller type guidance policy.

Action

Modify the planner input and try again.

HWNLM0507E The zone per fabric policy was violated.

Explanation

The Zone Planner cannot satisfy the zone per fabric guidance policy.

Action

Modify the planner input and try again.

HWNLM0508I Zone Set *zone set name* was created.

Explanation

None.

Action

None.

HWNLM0509I Zone *zone name* was created.

Explanation

None.

Action

None.

HWNLM0510I Zone *zone name* was added to Zone Set *zone set name*.

Explanation

None.

Action

None.

HWNLM0511I Host port *port id* was added to Zone *zone name*.

Explanation

None.

Action

None.

HWNLM0512I Subsystem port *port id* was added to Zone *zone name*.

Explanation

None.

Action

None.

HWNLM0513I SAN Planner started with guidance Policy (*guidance*), validation policies (*validation*), fabric WWN (*fabricWWN*), and using active zone set (*zone set name*).

Explanation

None.

Action

None.

HWNLM0514I The Zone Planner completed.

Explanation

None.

Action

None.

HWNLM0515W Not considering Subsystem {0} since Planner could not obtain its information from database (it could be undetectable).

Explanation

None.

Action

None.

HWNLMO516E Could not find any subsystem with the given input. The subsystem(s) may be undetectable.

Explanation

None.

Action

None.

HWNLMO517W Invalid candidate subsystem {0} for performance data for SAN Planner.

Explanation

None.

Action

None.

HWNLMO518I Not considering Pool: {0} of Subsystem: {1} since it does not have enough allocatable space.

Explanation

None.

Action

None.

HWNLMO519W Not considering Pool: {0} of Subsystem: {1} since it is not online.

Explanation

None.

Action

None.

HWNLMO520I Not considering Pool: {0} of Subsystem: {1} because it is not Fixed Block pool.

Explanation

None.

Action

None.

HWNLM0521I Not considering Pool: {0} of Subsystem: {1} because it is Solid State Disk pool.

Explanation

None.

Action

None.

HWNLM0522I Not considering Pool: {0} of Subsystem: {1} because it is not a Solid State Disk pool.

Explanation

None.

Action

None.

HWNLM0523I Not considering Pool: {0} of Subsystem: {1} because its Encryption Group is not matching with the input.

Explanation

None.

Action

None.

HWNLM0524I Not considering Pool: {0} of Subsystem: {1} because it is not thin-provisioning enabled.

Explanation

None.

Action

None.

HWNLM0525I Not considering Pool: {0} of Subsystem: {1} because its lock behavior is not matching with the input.

Explanation

None.

Action

None.

HWNLM0526I Not considering Pool: {0} of Subsystem: {1} because it is not in the selected input pools list or it is already filtered out.

Explanation

None.

Action

None.

HWNLM0527I Not considering Pool: {0} of Subsystem: {1} because its RAID level does not match with the input RAID level.

Explanation

None.

Action

None.

HWNLM0528I Considering volume: {0} of Subsystem: {1} as candidate since it is of the right size, is unassigned and not in a known replication relationship (or these requirements were over-ridden by user).

Explanation

None.

Action

None.

HWNLM0529W The subsystem {0} does not have performance data for the specified time interval.

Explanation

None.

Action

None.

HWNLM0530I Not considering volume {0} for new storage because the user indicated so.

Explanation

None.

Action

None.

HWNLM0531I Not considering the subsystem {0} because it does not support Extent Space Efficient volumes.

Explanation

None.

Action

None.

HWNLM0532I Not considering the Subsystem: {0} because it does not support Track Space Efficient volumes.

Explanation

None.

Action

None.

HWNLM0533I Not considering Pool: {0} of Subsystem: {1} because it does not have repository capacity available/defined for Track Space Efficient volumes.

Explanation

None.

Action

None.

HWNLM0534I New Capacity Planning Advice Task Started....

Explanation

None.

Action

None.

HWNLM0535I Capacity Planning Advice Task Completed.

Explanation

None.

Action

None.

HWNL0536I Not considering Pool: {0} of Subsystem: {1} because its backend storage RAID level does not match the input RAID level

Explanation

None.

Action

None.

HWNL0537I Not considering Pool: {0} of Subsystem: {1} because its backend storage is not configured into this IBM Spectrum Control or is currently undetectable and input requires a specific backend RAID level

Explanation

None.

Action

None.

HWNL0538I Not considering IO Group: {0} of Subsystem: {1} because it does not have two nodes associated with it

Explanation

None.

Action

None.

HWNL0539I Not considering IO Group: {0} of Subsystem: {1} because it has more than or equal to 2048 vdisks

Explanation

None.

Action

None.

HWNL0540I Not considering IO Group: {0} of Subsystem: {1} because it does not have enough available mirroring memory

Explanation

None.

Action

None.

HWNLM0541I Not considering Volume: {0} of Subsystem: {1} because its subsystem was found to be unacceptable based on inputs

Explanation

None.

Action

None.

HWNLM0542I Not considering Volume: {0} of Subsystem: {1} because its pool was found to be unacceptable based on inputs

Explanation

None.

Action

None.

HWNLM0543I Not considering Volume: {0} of Subsystem: {1} because its IO group was found to be unacceptable based on inputs

Explanation

None.

Action

None.

HWNLM0544I Not considering Volume: {0} of Subsystem: {1} because its size does not match the input

Explanation

None.

Action

None.

HWNLM0545I Not considering Subsystem: {0} because not all of its backend subsystems are configured into this IBM Spectrum Control or some are currently undetectable. This is unacceptable when an input RAID level is specified or a workload profile other than space-only is selected

Explanation

None.

Action

None.

HWNLMO546I Not considering Subsystem: {0} because IBM Spectrum Control does not have adequate performance data for it and/or some of its backend subsystems. For planning with workload profiles other than space-only, this subsystem and all of its backend subsystems need to have daily performance data

Explanation

None.

Action

None.

HWNLMO547I Not considering Subsystem: {0} since none of its IO groups have enough available mirroring memory to support vdisk mirroring input

Explanation

None.

Action

None.

HWNLMO548I Not considering Subsystem: {0} since it does not support vdisk mirroring (code level is below v4.3)

Explanation

None.

Action

None.

HWNLMO549I Not considering Subsystem: {0} since it does not support space-efficient vdisks (code level is below v4.3)

Explanation

None.

Action

None.

HWNLMO550I Not considering IO Group: {0} of Subsystem: {1} because it does not appropriate connectivity to hosts selected in the

input

Explanation

None.

Action

None.

HWNLM0551I Not considering Subsystem: {0} since none of its IO groups have appropriate connectivity to hosts selected in the input

Explanation

None.

Action

None.

HWNLM0552I Not considering Volume: {0} of Subsystem {1} since it does not match input on thin-provisioning characteristics

Explanation

None.

Action

None.

HWNLM0553I Volume size too small {0} MB. Minimum size should be {1} MB for the pool {2}. Not considering this pool.

Explanation

None.

Action

None.

HWNLM0554I Volume size too big {0} MB. Maximum size should be {1} MB for the pool {2}. Not considering this pool.

Explanation

None.

Action

None.

HWNLMO555I Not considering Volume: {0} of Subsystem: {1} because it is in a known replication relationship.

Explanation

None.

Action

None.

HWNLMO556I Volume {0} can not be moved to the following candidate pool(s) due to insufficient allocatable space: {1}.

Explanation

None.

Action

None.

HWNLMO557I The max theoretical I/O capability of storage pool {0} on subsystem {1} has not been set. Using default value {2}.

Explanation

Because insufficient information was available to calculate the maximum I/O capability of the pool, a default value was set.

Action

If the I/O capability of the storage pool cannot be determined, the user-defined value for theoretical I/O capability is used. If this value has not been set, the default value is used and the informational message HWNLMO557I is generated. To add the information, go to the Pools page, right-click the pool, and then click View Properties. On the Back-End Storage tab, click Edit.

HWNLMO558W The subsystem pool {0} does not have performance data for the specified time interval.

Explanation

None.

Action

None.

HWNLMO559I Not considering IO Group: {0} of Subsystem: {1} because compression was specified, but the IO Group does not have compression active.

Explanation

Compression was specified, but the IO group could not satisfy the request because there are no active compressed volumes in the IO group.

Action

If no IO groups have compression active, then create a compressed volume in the IO group to make compression active.

HWNLM0560I Not considering Volume: {0} of Subsystem {1} since it is not compressed when compression was specified.

Explanation

Compression was specified, but the volume could not satisfy the request because it is not compressed.

Action

Select volumes that are compressed.

HWNLM0560W Fabric agent not available to perform the zoning operation on the fabric {0}.

Explanation

None.

Action

None.

HWNLM0561I Fabric agent available to perform the zoning operation on the fabric {0}.

Explanation

None.

Action

None.

HWNLM0562E Fabric service exception occurred trying to check for zone control on the fabric {0}.

Explanation

None.

Action

None.

HWNLM0563W A zone *unnecessary_zone* was not created, because the host already has connectivity to the storage system using *existing_zone*.

Explanation

When automatic zoning is enabled, new zones might be created during provisioning to connect the server or hypervisor to the storage system. New zones are created only if the server or hypervisor does not have connectivity with the storage system. In this case, a new zone was not needed because another zone between the same ports already exists in the active zone set.

Action

No action is required.

HWNLM0564W There is already a zone named *zone_name*. A zone named *new_zone_name* was created instead.

Explanation

The planned zone name was not used because a zone with that name exists. Although the zone was created with a different name, the provisioning task shows the original planned name.

Action

No action is required.

HWNLM0565W There is already a volume named *vol_name*. A volume named *new_vol_name* will be created instead.

Explanation

The planned volume name was not used because a volume with that name exists. Although the volume will be created with a different name, the provisioning task shows the original planned name.

Action

No action is required.

HWNLM0566I Not considering Pool: {0} of Subsystem: {1} because it is mixed pool.

Explanation

None.

Action

None.

HWNLM0567I Not considering Subsystem: {0} for vdisk mirroring since it does not have at least two acceptable mdiskgroups from different backend subsystems.

Explanation

None.

Action

None.

HWNLM0568I Not considering I/O group {0} of subsystem {1} because it already reached the limit of 200 compressed vdisks.

Explanation

An I/O group can contain a maximum of 200 compressed vdisks, and the I/O group already reached this limit.

Action

No action is required.

HWNLM0569I Not considering subsystem {0} because none of its I/O groups can be used for compressed vdisks.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNLM0570I Not considering I/O group {0} of subsystem {1} because host definition {2} has access restrictions.

Explanation

The host definition has access restrictions and can can not connect to the I/O group.

Action

No action is required.

HWNLM0571I Not considering storage system {0} because none of its I/O groups allow restricted access to all hosts.

Explanation

The host definition has access restrictions and can can not connect to any of the I/O groups of the storage system.

Action

No action is required.

HWNLM801I Synchronous Refresh Configuration of the Storage Subsystem was completed Successfully for Subsystem *subsystem*.

Explanation

None.

Action

None.

HWNLM802E Synchronous Refresh Configuration of the Storage Subsystem did not complete successfully for Subsystem *subsystem*.

Explanation

None.

Action

None.

HWNL803I Replication Session was created successfully *session*.

Explanation

None.

Action

None.

HWNL804E Replication Session creation failed. ReplicationManager [*session*].

Explanation

None.

Action

None.

HWNL805I Successful check for existence of session *session*.

Explanation

None.

Action

None.

HWNL806E No Such session exists with name *session*.

Explanation

None.

Action

None.

HWNL807I CopySets Added to Session successfully *session*.

Explanation

None.

Action

None.

HWNL808E CopySets addition to session failed. ReplicationManager [*session*].

Explanation

None.

Action

None.

HWNLM809E CopySets creation failed. ReplicationManager [session].

Explanation

None.

Action

None.

HWNLM810I Storage Subsystem Configuration refreshed successfully in Replication Manager ss.

Explanation

None.

Action

None.

HWNLM811E Storage Subsystem Configuration refresh operation failed in Replication Manager [ss].

Explanation

None.

Action

None.

HWNLM812E Fabric agent not available to perform the zoning operation on the fabric {0}.

Explanation

None.

Action

None.

HWNLM813I Replication Session was started successfully session.

Explanation

None.

Action

None.

HWNL814W Replication Session start failed. ReplicationManager [session] .

Explanation

None.

Action

None.

HWNL815W Not considering the Storage Subsystem {0} because it is not registered in TPC-R.

Explanation

None.

Action

None.

HWNL816I Please use Replication Manager console to schedule point-in-time copy creation.

Explanation

None.

Action

None.

HWNL817W No cluster partnership exists between source subsystem [srcss] and target subsystem [tgtss].

Explanation

Cluster partnership needs to established between source and target storage subsystems

Action

Cluster partnership needs to established between source and target storage subsystems

HWNL818W No path exists between source subsystem [srcss] and target subsystem [tgtss].

Explanation

Connectivity path needs to established between source and target storage subsystems

Action

Connectivity path needs to established between source and target storage subsystems

HWNLM820I Copysets can not be added to a session when it's in Prepared/Suspend/Recover/Flash state. Please try again later after some time. State of session can be checked from Replication Manager console.

Explanation

Copysets can not be added to a session when it's in Prepared/Suspend/Recover/Flash state. Please try again later after some time. State of session can be checked from Replication Manager console.

Action

Copysets can not be added to a session when it's in Prepared/Suspend/Recover/Flash state. Please try again later after some time. State of session can be checked from Replication Manager console.

HWNLM0821W IBM Spectrum Control had no information about the Fibre Channel configuration, so it did not verify fabric connectivity or change the zoning configuration.

Explanation

Because IBM Spectrum Control did not have any information about the Fibre Channel configuration, it did not verify fabric connectivity and did not change the zoning configuration.

Action

After the provisioning operation completes, ensure that the involved storage subsystems and hosts are connected through Fibre Channel and create and activate the necessary zones and zone sets in order to make the connection possible.

HWNOP - Storage optimizer messages

- [HWNOP0001I The Collection phase of the integrated Storage Optimizer job jobname has started.](#)
- [HWNOP0002E The Collection phase of the integrated Storage Optimizer job jobname has failed with errors. Message from exception: message.](#)
- [HWNOP0003I The Collection phase of the integrated Storage Optimizer job jobname has completed.](#)
- [HWNOP0004W The Collection phase of the integrated Storage Optimizer job jobname has completed with warnings.](#)
- [HWNOP0005E The Collection phase of the integrated Storage Optimizer job jobname has completed with errors.](#)
- [HWNOP0006I Zone set zoneset created on fabric fabric.](#)
- [HWNOP0007I The Collection phase has started Collecion for the Analysis job jobname, on Subsystem: subsystem, Start Date: startdate, End Date: enddate](#)
- [HWNOP0008I The Collection phase has started Collection for the Optimization job creator, Analysis Job ID: analysisjob, on Subsystem: subsystem](#)
- [HWNOP0009W The subsystem subsystem does not exist in the internal database.](#)
- [HWNOP0010W The Performance data summary level is unknown for subsystem subsystem.](#)
- [HWNOP0011I Collecting Performance Data for Subsystem subsystem Storage Pool pool with summary level level for time range range](#)
- [HWNOP0012E Failed to bind to Performance Manager Service](#)
- [HWNOP0013E Database operation failed error](#)
- [HWNOP0014I Optimizer Service Started](#)
- [HWNOP0015I Optimizer Collector called for Job job with inputs input](#)
- [HWNOP0016I Collecting Configuration Data for Subsystem subsystem](#)
- [HWNOP0017I Updating Aggregated Workload Profile Data Table in Database for Subsystem subsystem](#)
- [HWNOP0018I The collection of configuration and performance data is finished. Now starting Analysis...](#)
- [HWNOP0019I The Analyzer is computing utilizations for storage subsystem subsystem.](#)
- [HWNOP0020I The Analyzer is updating the database with utilizations for storage subsystem subsystem.](#)
- [HWNOP0021I The Analyzer is reading analysis data from the database.](#)
- [HWNOP0022I The Analyzer queued the job on the planner.](#)
- [HWNOP0023I The Analyzer is done for the job.](#)
- [HWNOP0024I The Planner is starting the planning operation.](#)
- [HWNOP0025I The Planner generated number recommendations](#)
- [HWNOP0026I The Planner queued the job on the effector.](#)
- [HWNOP0027I The Effector is adding the plans in the database.](#)
- [HWNOP0028I The Effector is done. The Optimization job is complete.](#)
- [HWNOP0029I Starting Planning phase.](#)
- [HWNOP0030I Starting Consolidation Plan.](#)
- [HWNOP0031I The Storage Optimizer cannot migrate or consolidate volume volume from source pool pool on subsystem subsystem because of space or utilization threshold constraints.](#)
- [HWNOP0032I The Optimizer consolidated volume volume to Storage Pool pool. The Storage Pool size is size. The Max Size is maxsize.](#)
- [HWNOP0033I There were duration day\(s\) in the selected interval for subsystem subsystem. numData of these had performance data available.](#)

- [HWNOP0034E No data performance data is available for the subsystem subsystem and the time interval that you have selected. Please choose a new start and end date interval to analyze.](#)
- [HWNOP0035I volume volumes out of maxvolume volumes were not consolidated.](#)
- [HWNOP0036E The Optimizer job failed with the following error: error](#)
- [HWNOP0037I Performance Data for Subsystem subsystem Storage Pool pool will not be collected as it is of mixed pool type.](#)
- [HWNOP0038I The Storage Optimizer job has been queued for processing.](#)

HWNOP0001I The Collection phase of the integrated Storage Optimizer job *jobname* has started.

Explanation

This message is for informational purposes. No further investigation is required.

Action

No action required.

HWNOP0002E The Collection phase of the integrated Storage Optimizer job *jobname* has failed with errors. Message from exception: *message*.

Explanation

The Storage Optimizer job failed due to an unexpected error during the data collection phase. Look at the log files for more information.

Action

Make sure the IBM Spectrum Control database is online and run the job again. If the error still persists, Contact IBM customer technical support with all related errors.

HWNOP0003I The Collection phase of the integrated Storage Optimizer job *jobname* has completed.

Explanation

None.

Action

No Action Required.

HWNOP0004W The Collection phase of the integrated Storage Optimizer job *jobname* has completed with warnings.

Explanation

None.

Action

No Action Required.

HWNOP0005E The Collection phase of the integrated Storage Optimizer job *jobname* has completed with errors.

Explanation

The Storage Optimizer job failed due to an unexpected error. Look at the log files for more information.

Action

Make sure the IBM Spectrum Control database is online and run the job again. If the error still persists, Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNOP0006I Zone set *zoneset* created on fabric *fabric* .

Explanation

None.

Action

No Action Required.

HWNOP0007I The Collection phase has started Collecion for the Analysis job *jobname*, on Subsystem: *subsystem* , Start Date: *startdate* , End Date: *enddate*

Explanation

None.

Action

No Action Required.

HWNOP0008I The Collection phase has started Collection for the Optimization job *creator*, Analysis Job ID: *analysisjob* , on Subsystem: *subsystem*

Explanation

None.

Action

No Action Required.

HWNOP0009W The subsystem *subsystem* does not exist in the internal database .

Explanation

None.

Action

No Action Required.

HWNOP0010W The Performance data summary level is unknown for subsystem *subsystem* .

Explanation

None.

Action

No Action Required.

HWNOP0011I Collecting Performance Data for Subsystem *subsystem*
Storage Pool *pool* with summary level *level* for time range *range*

Explanation

None.

Action

No Action Required.

HWNOP0012E Failed to bind to Performance Manager Service

Explanation

The Storage Optimizer job failed due to an unexpected error. Look at the log files for more information.

Action

Make sure the IBM Spectrum Control database is online and run the job again. If the error still persists, Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNOP0013E Database operation failed *error*

Explanation

The Storage Optimizer job failed due to an unexpected error. Look at the log files for more information.

Action

Make sure the IBM Spectrum Control database is online and run the job again. If the error still persists, Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNOP0014I Optimizer Service Started

Explanation

None.

Action

No Action Required.

HWNOP0015I Optimizer Collector called for Job *job* with inputs *input*

Explanation

None.

Action

No Action Required.

HWNOP0016I Collecting Configuration Data for Subsystem *subsystem*

Explanation

None.

Action

No Action Required.

HWNOP0017I Updating Aggregated Workload Profile Data Table in Database for Subsystem *subsystem*

Explanation

None.

Action

No Action Required.

HWNOP0018I The collection of configuration and performance data is finished. Now starting Analysis...

Explanation

None.

Action

No Action Required.

HWNOP0019I The Analyzer is computing utilizations for storage subsystem *subsystem*.

Explanation

None.

Action

No Action Required.

HWNOP0020I The Analyzer is updating the database with utilizations for storage subsystem *subsystem*.

Explanation

None.

Action

No Action Required.

HWNOP0021I The Analyzer is reading analysis data from the database.

Explanation

None.

Action

No Action Required.

HWNOP0022I The Analyzer queued the job on the planner.

Explanation

None.

Action

No Action Required.

HWNOP0023I The Analyzer is done for the job.

Explanation

None.

Action

No Action Required.The Optimizer has completed analysis of the collected data and performance projections.

HWNOP0024I The Planner is starting the planning operation.

Explanation

None.

Action

No Action Required.

HWNOP0025I The Planner generated *number* recommendations

Explanation

None.

Action

No Action Required.

HWNOP0026I The Planner queued the job on the effector.

Explanation

None.

Action

No Action Required.

HWNOP0027I The Effector is adding the plans in the database.

Explanation

None.

Action

No Action Required.

HWNOP0028I The Effector is done. The Optimization job is complete.

Explanation

None.

Action

No Action Required.

HWNOP0029I Starting Planning phase.

Explanation

None.

Action

No Action Required.

HWNOP0030I Starting Consolidation Plan.

Explanation

None.

Action

No Action Required.

HWNOP0031I The Storage Optimizer cannot migrate or consolidate volume *volume* from source pool *pool* on subsystem *subsystem* because of space or utilization threshold constraints.

Explanation

None.

Action

No Action Required.

HWNOP0032I The Optimizer consolidated volume *volume* to Storage Pool *pool*. The Storage Pool size is *size*. The Max Size is *maxsize*.

Explanation

None.

Action

No Action Required.

HWNOP0033I There were *duration* day(s) in the selected interval for subsystem *subsystem*. *numData* of these had performance data available.

Explanation

None.

Action

No Action Required.

HWNOP0034E No *data* performance data is available for the subsystem *subsystem* and the time interval that you have selected. Please choose a new start and end date interval to analyze.

Explanation

None

Action

Make sure to run performance monitors to collect performance data.

HWNOP0035I *volume* volumes out of *maxvolume* volumes were not consolidated.

Explanation

None

Action

Make sure to run performance monitors to collect performance data.

HWNOP0036E The Optimizer job failed with the following error: *error*

Explanation

The Storage Optimizer job failed due to an unexpected error. Look at the log files for more information.

Action

None

HWNOP0037I Performance Data for Subsystem *subsystem* Storage Pool *pool* will not be collected as it is of mixed pool type.

Explanation

None.

Action

No Action Required.

HWNOP0038I The Storage Optimizer job has been queued for processing.

Explanation

The Storage Optimizer job has been queued for processing.

Action

No action is required.

HWNPM - Performance manager messages

- [HWNPM0001E The specified summarization level \(level\) is invalid. It must be an integer value between minimum and maximum, inclusive.](#)
- [HWNPM0002E The specified device category \(category\) is invalid. It must be an integer value between minimum and maximum, inclusive.](#)
- [HWNPM0003E The specified device type \(type\) is invalid. It must be an integer value between minimum and maximum, inclusive.](#)
- [HWNPM0004E The specified component type \(type\) is invalid. It must be an integer value between minimum and maximum, inclusive.](#)
- [HWNPM0006E The string specified as parameter \(string\) exceeded its allowed length \(maximum length\).](#)
- [HWNPM0007E The value specified as parameter \(value\) is invalid.](#)
- [HWNPM0008E A required parameter is missing \(null\).](#)
- [HWNPM0010E The specified device ID \(device ID\) is invalid. It must conform to the pattern 'name+nameFormat'.](#)
- [HWNPM0011E The specified component ID \(component ID\) is invalid. It must be a simple WWN \(16 hexadecimal characters\).](#)
- [HWNPM0012E The specified component ID \(component ID\) was not found or is not unique in the IBM Spectrum Control database.](#)
- [HWNPM0013E The specified component ID \(component ID\) is invalid.](#)
- [HWNPM0015E Failed to retrieve the requested data because the service is unavailable.](#)
- [HWNPM0021E The device identifier specified as parameter \(device ID\) is invalid.](#)
- [HWNPM0090E Failed to retrieve the requested data because the service is unavailable.](#)
- [HWNPM0099E The requested operation failed because of an internal error.](#)
- [HWNPM0101E Unable to create the specified performance service instance \({0}\).](#)
- [HWNPM0200I This operation \(operation name\) on Performance Manager was successful.](#)
- [HWNPM0201E The device \(device_id\) that was passed to the method is invalid.](#)
- [HWNPM0202E The device category \(device_category\) that was passed to the method is invalid.](#)
- [HWNPM0203E The device type received \(device_type\) is invalid.](#)
- [HWNPM0204E The device type - HOST - that was passed to the method is not supported.](#)
- [HWNPM0205E The specified performance collection policy is invalid.](#)
- [HWNPM0209I The device type and device category are valid.](#)
- [HWNPM0210E Collector failed to start due to system failure.](#)
- [HWNPM0220E Collector failed to stop due to system failure.](#)
- [HWNPM0230E One or more of the specified performance collection policies are invalid.](#)

- [HWNPM0231W The specified performance collection policy is ignored because it conflicts with another policy in the same parameter list.](#)
- [HWNPM0232E The specified performance collection policy contains an unsupported interval length.](#)
- [HWNPM0233E The specified performance collection policy contains an unsupported frequency.](#)
- [HWNPM0234E The specified performance collection policy contains an unsupported duration.](#)
- [HWNPM0240E The attempt to update the specified performance collection policies has failed.](#)
- [HWNPM0241E The attempt to reset the specified performance collection policies has failed.](#)
- [HWNPM0242E The attempt to remove the specified performance collection policies has failed.](#)
- [HWNPM0249W An attempt to dynamically update one or more running performance collectors with a new performance collection policy has failed.](#)
- [HWNPM0250E One or more default performance collection policies are missing from the database.](#)
- [HWNPM0281I Performance monitoring is unavailable for resource resource_name because an agent for monitoring the resource was not defined. For IBM Spectrum Scale, the problem might occur because the data collection service cannot connect to port 9084 on the node where the collector component of the IBM Spectrum Scale performance monitoring tool is running.](#)
- [HWNPM0282I Performance monitoring is unavailable for resource resource_name because the associated data sources are unable to collect performance data from the resource.](#)
- [HWNPM0283I Performance monitoring is unavailable for resource_name because this resource does not support performance monitoring.](#)
- [HWNPM0284I Performance monitoring is unavailable for resource resource_name because the associated agent does not have the required level of software agent_level.](#)
- [HWNPM0285I Performance monitoring is unavailable for resource resource_name because the associated agent is unable to fully monitor the resource.](#)
- [HWNPM0286I Performance monitoring is unavailable for resource resource_name because the associated SMI-S provider does not have the required SMI-S support.](#)
- [HWNPM0287I Performance monitoring is unavailable for resource resource_name because the resource or the associated agent does not support performance monitoring.](#)
- [HWNPM0288I Performance monitoring is unavailable for resource resource_name because the resource was not probed.](#)
- [HWNPM0289W Performance monitoring is unavailable for resource resource_name because no agents are available.](#)
- [HWNPM0290E Performance monitoring is unavailable for resource resource_name because the associated agent was could not be selected.](#)
- [HWNPM0291I Performance monitoring is unavailable for switch resource_name because the switch has no ports.](#)
- [HWNPM0292I Performance monitoring is unavailable for switch resource_name because the switch was not probed using the correct agent.](#)
- [HWNPM0293I Performance monitoring is unavailable for FlashSystem resource_name because its SNMP agent is disabled. You can enable SNMP for FlashSystem storage systems in the FlashSystem GUI<a>.](#)
- [HWNPM0300E There is an exception for each device processed in a multiple devices call.](#)
- [HWNPM0390E A system failure occurred.](#)
- [HWNPM0400I This operation \(operation_name\)on Threshold Service was successful.](#)
- [HWNPM0401E The device \(device_id\) that was passed to the method is invalid.](#)
- [HWNPM0410E The Performance threshold policy that was passed to the method \(threshold_policy\)is null.](#)
- [HWNPM0411E The Performance threshold that was passed to the method \(threshold\)is null.](#)
- [HWNPM0412E The Performance threshold filter that was passed to the method \(filter\)is null.](#)
- [HWNPM0420E The device type received \(device_type\) is invalid.](#)
- [HWNPM0421E There is no default performance threshold policy or default threshold filter for this device.](#)
- [HWNPM0425E There is an exception for each device processed in a multiple devices threshold call.](#)
- [HWNPM0590E Performance Manager failed due to system failure.](#)
- [HWNPM0600E Parameter number a number of the call made to the IBM Spectrum Control Performance Manager reporting API method name of the api is invalid. The invalid value is the invalid value of parameter descriptive name of the parameter.](#)
- [HWNPM0601E A request to continue the data retrieval can not be performed. Information for continuing the data retrieval does not exist.](#)
- [HWNPM0602E Support for the device type device_type name is not available in the Performance Manager reporting API function method name.](#)
- [HWNPM0603E The performance reporting API method_name failed at time_of_failure as a result of an internal processing exception. The Performance Manager logs contain message message_ID that describes the internal processing exception.](#)
- [HWNPM0604E The sort order parameter of a call to the performance reporting API method_name contains a value not included in the report columns list, at position list_item in the sort order list.](#)
- [HWNPM0605E Performance metadata cannot be displayed.](#)
- [HWNPM0606E Unable to instantiate performance reporting service service_class_name.](#)
- [HWNPM0607E An error occurred while the performance data was being retrieved.](#)
- [HWNPM0630E An invalid operator \(operator identifier\) was specified for the filter expression.](#)
- [HWNPM0631E An invalid first operand \(operand class\) was specified for the filter expression. It must be a operand class class.](#)
- [HWNPM0632E An invalid first operand was specified for the filter expression. The data type of the operand \(data type\) is invalid or unsupported.](#)
- [HWNPM0633E An invalid second operand \(operand class\) was specified for the filter expression. It must be a operand class class.](#)
- [HWNPM0650E The IBM Spectrum Control Performance Manager reporting API method name failed as a result of exception the related exception from a call to method method_name, of the lower-level service name internal service.](#)
- [HWNPM0651E The configuration data needed to generate the affected volumes and hosts report for the device device_name was not found in the IBM Spectrum Control database.](#)
- [HWNPM0652E The requested performance metrics are no longer available. Clear your browser cache to proceed.](#)
- [HWNPM2000I Performance monitoring is enabled.](#)
- [HWNPM2001E The IBM Spectrum Control Performance Manager is not operational.](#)
- [HWNPM2002E An initialization error occurred.](#)
- [HWNPM2003E Initialization of the Device server event service failed. No performance threshold exception alerts will be generated.](#)
- [HWNPM2004E Initialization of the product scheduler status service failed. The status of performance monitors will not be updated in the GUI.](#)
- [HWNPM2005E Initialization of the product configuration data service failed. Performance monitors cannot be started without this service.](#)
- [HWNPM2006E Initialization of the product configuration data service failed. Performance monitors cannot be started without this service.](#)
- [HWNPM2007E Initialization of the product counter data service failed for device type using agent type. Performance monitors will not be able to collect performance data from devices of this type.](#)
- [HWNPM2008E Initialization of the product metadata service failed. Performance monitors cannot be started without this service.](#)
- [HWNPM2009E Unable to instantiate lower level service service_class_name.](#)
- [HWNPM2010E Unable to instantiate the collection logic implementation service class name.](#)
- [HWNPM2011E Unable to instantiate the performance statistics data class class_name.](#)
- [HWNPM2012I The product is using trace log directory log_directory_name.](#)
- [HWNPM2020W The performance monitor for device device_name is not currently active, so a dynamic update of its monitor policy is not necessary.](#)
- [HWNPM2021W The performance monitor for device device_name is not currently active, so a dynamic update of its threshold policy is not necessary.](#)
- [HWNPM2022E A performance monitor for device device_name is already active. A new monitor for the same device cannot be started until the previous monitor completes or is cancelled.](#)
- [HWNPM2023W The performance monitor for device device_name is not currently active.](#)
- [HWNPM2024E Unable to find a monitor policy applicable to resource resource_name.](#)

- [HWNPM2025E Unable to find a threshold policy applicable to resource resource_name.](#)
- [HWNPM2026I The performance monitor's primary process has failed unexpectedly. Attempting to recover from the failure.](#)
- [HWNPM2027I The performance monitor threshold checker has failed unexpectedly. Attempting to recover from the failure.](#)
- [HWNPM2028I The performance monitor purge process has failed unexpectedly. Attempting to recover from the failure.](#)
- [HWNPM2029I Successfully recovered from the performance monitor failure.](#)
- [HWNPM2030E Unable to recover from the performance monitor failure. The performance monitor for the storage resource will be shut down.](#)
- [HWNPM2031E The performance monitor failed due to an internal error.](#)
- [HWNPM2032W The performance monitor for device device_name is not currently using the default monitor policy, so a dynamic update of the policy is not necessary.](#)
- [HWNPM2033W The performance monitor for device device_name is not currently using the default threshold policy, so a dynamic update of the policy is not necessary.](#)
- [HWNPM2040E The device key specified for the snapshot vote \(key\) was not found in the database. The device does not exist.](#)
- [HWNPM2050E Failed to get the latest configuration data for device device_name.](#)
- [HWNPM2051E No performance data was collected from device device_name for the current collection interval \(time_stamp\) because the performance monitor was stopped.](#)
- [HWNPM2052E No performance data was collected from device device_name for the current collection interval due to an error. Data was last collected at time_stamp.](#)
- [HWNPM2053E The new performance data collected from device device_name could not be saved in the database. Increase the size of the transaction log.](#)
- [HWNPM2054E The new performance data collected from device device_name could not be saved in the database. Increase the size of the database lock list.](#)
- [HWNPM2055E The new performance data collected from resource device_name could not be saved.](#)
- [HWNPM2056E No performance data was collected from device device_name for the current performance monitor job duration. The performance monitor job status is set to 'failed'.](#)
- [HWNPM2057E No performance data was collected from device device_name for the current collection interval because the performance monitor was stopped.](#)
- [HWNPM2058E No performance data was collected from device device_name for the current collection interval due to an error.](#)
- [HWNPM2060W The device does not support performance management for segment pool pool_ID. Only incomplete performance data can be collected for array array_ID.](#)
- [HWNPM2061W The device does not support performance management for segment pool pool_ID. Only incomplete performance data can be collected for device adapter DA_ID.](#)
- [HWNPM2062W Invalid error message saved in database](#)
- [HWNPM2100E The performance monitor for resource device_name cannot be started because configuration data for the resource is not available.](#)
- [HWNPM2101E All agents that can collect performance data for resource device_name are currently non-operational.](#)
- [HWNPM2102E The performance monitor for resource device_name cannot be started because the resource might not support the collection of performance data.](#)
- [HWNPM2103W Agent agent_name is non-operational. Attempting to find an alternative agent.](#)
- [HWNPM2104I The performance monitor policy was adjusted due to agent limitations. Current values in effect are: interval=length in seconds secs, frequency=length in seconds secs.](#)
- [HWNPM2105E The performance monitor for resource resource_name failed because the resource for enabling performance data collection cannot be reached.](#)
- [HWNPM2106E The performance monitor for device device_name failed because of errors trying to enable performance data collection on the device or device agent: error description](#)
- [HWNPM2107E The performance monitor for device device_name failed because of unrecognized errors trying to enable performance data collection on the device or device agent: error description](#)
- [HWNPM2108E The performance monitor for resource resource_name failed during shutdown because the resource cannot be reached for terminating data collection.](#)
- [HWNPM2109E The performance monitor for resource resource_name failed during shutdown because of errors during termination of performance data collection: error description](#)
- [HWNPM2110E The performance monitor for resource resource_name failed during shutdown because of unrecognized errors during termination of performance data collection: error description](#)
- [HWNPM2111E The performance monitor for resource resource_name failed because of errors retrieving the most recent configuration data for the resource.](#)
- [HWNPM2112I Agent agent_name was selected for performance data collection from resource resource_name.](#)
- [HWNPM2113I The performance monitor for resource resource_name is starting in an active state.](#)
- [HWNPM2114I The performance monitor for resource resource_name is starting in a dormant state.](#)
- [HWNPM2115I Monitor Policy: name="policy_name", creator="policy_creator", description="policy_description"](#)
- [HWNPM2116I Monitor Policy: retention period: sample data=length in days days, hourly data=length in days days, daily data=length in days days.](#)
- [HWNPM2117I Monitor Policy: interval length=length in seconds secs, frequency=length in seconds secs, duration=length in hours hours.](#)
- [HWNPM2118I Threshold Policy: name="policy_name", creator="policy_creator", description="policy_description"](#)
- [HWNPM2119I Threshold Policy: retention period: exception data=length in days days.](#)
- [HWNPM2120I Threshold Policy: threshold name=name, component=component type, enabled=Yes or No, boundaries=critical stress boundary, warning stress boundary, warning idle boundary, critical idle boundary units.](#)
- [HWNPM2121I Monitor Policy: interval length=length in seconds secs, frequency=length in seconds secs, duration=continue indefinitely.](#)
- [HWNPM2122W No valid performance data was provided by the monitored resource. No performance data records were inserted into the database.](#)
- [HWNPM2123I Performance data for resource timestamp date and time was collected and processed successfully, record count performance data records were inserted into the database repository.](#)
- [HWNPM2124W Performance data continuity is broken. The device was possibly reset or rebooted, record count performance data records were discarded.](#)
- [HWNPM2125W Aggregated performance values have been computed from the remaining data records, but their accuracy cannot be guaranteed.](#)
- [HWNPM2126I The performance monitor for device device_name is stopping because its intended duration has elapsed.](#)
- [HWNPM2127I The performance monitor for device device_name is stopping due to a user request.](#)
- [HWNPM2128E The performance monitor for device device_name is stopping due to an unexpected failure.](#)
- [HWNPM2129I The performance monitor for device device_name is stopping because of a shutdown request.](#)
- [HWNPM2130W Failed to retrieve the latest configuration data for device device_name.](#)
- [HWNPM2131W Performance data could not be collected for device device_name, because the device or data source cannot be reached \(reason reason code\). The current samples are skipped.](#)
- [HWNPM2132W Performance data could not be collected for device device_name. The current samples are skipped. \(error description\)](#)
- [HWNPM2133W Performance data could not be collected for device device_name due to an unknown error. The current samples are skipped.](#)
- [HWNPM2134W The state of the performance monitor for resource resource_name started, but the status of the performance monitor was not updated.](#)
- [HWNPM2135W The state of the performance monitor for device device_name has changed to 'active', but could not be recorded appropriately.](#)
- [HWNPM2136W The performance monitor for the resource resource_name generated a warning, but the status of the performance monitor was not updated.](#)
- [HWNPM2137W The performance monitor for the resource resource_name stopped, but the status of the performance monitor was not updated.](#)

- [HWNPM2138W The performance monitor for the resource resource name completed the collection of data, but the status of the performance monitor was not updated.](#)
- [HWNPM2139W The performance monitor for the resource resource name failed, but the status of the performance monitor was not updated.](#)
- [HWNPM2140W The status of the performance monitor for the resource resource name was not updated.](#)
- [HWNPM2141E The service is unavailable because an unexpected error occurred.](#)
- [HWNPM2142E Performance data can't be collected for the resource resource name because the performance monitor was disabled.](#)
- [HWNPM2143E The performance monitor for the resource resource name was started, but the status of the performance monitor was not updated and might not be shown in the GUI.](#)
- [HWNPM2144W The performance data cannot be checked against the alert conditions, so no alerts can be generated.](#)
- [HWNPM2145I The data is being collected by the data collector: data collector host.](#)
- [HWNPM2146W Performance data could not be collected for device device name, the exact reason for the failure could not be determined. The current samples are skipped.](#)
- [HWNPM2147W Performance data could not be collected for device device name, because of a bad target \(device or agent\) address. The current samples are skipped.](#)
- [HWNPM2148W Performance data could not be collected for device device name, because of an unknown target address. The current samples are skipped.](#)
- [HWNPM2149W Performance data could not be collected for device device name, because of an unreachable target address. The current samples are skipped.](#)
- [HWNPM2150W Performance data could not be collected for device device name, because of an unresponsive target. The current samples are skipped.](#)
- [HWNPM2151W Performance data could not be collected for device device name, because a communication time-out for communication that uses UDP rather than TCP. The current samples are skipped.](#)
- [HWNPM2200I The performance monitor successfully collected the configuration data for the storage system with the following internal resources: number of pools pools, number of controllers controllers, number of device adapters device adapters, number of ports ports, number of host connections host connections, number of ranks ranks, number of arrays arrays, and number of volumes volumes.](#)
- [HWNPM2201I The performance monitor successfully collected the configuration data for the storage system with the following internal resources: number of io groups I/O Groups, number of nodes nodes, number of ports ports, number of host connections host connections, number of pools pools, number of managed disks managed disks, number of local disks local disks, number of volumes volumes, and number of volume copies volume copies.](#)
- [HWNPM2202I The performance monitor successfully retrieved the configuration data for the switch. The following internal resources were found: number of trunks trunks, and number of ports ports.](#)
- [HWNPM2203I The performance monitor successfully retrieved the configuration data for the storage system. The following internal resources were found: number of host connections host connections, number of modules modules, number of ports ports, number of pools pools, and number of volumes volumes.](#)
- [HWNPM2204I The performance monitor successfully retrieved the configuration data for the storage system. The following internal resources were found: number of nodes nodes, number of ports ports, and number of modules flash modules.](#)
- [HWNPM2205I The performance monitor successfully retrieved the configuration data for the storage system. The following internal resources were found: number of ports ports, number of controllers controllers, number of volumes volumes, and number of disks disks.](#)
- [HWNPM3000E There was a problem establishing the database connection.](#)
- [HWNPM3001E An unexpected null row was returned from a database cursor.](#)
- [HWNPM3002E An unexpected database exception occurred.](#)
- [HWNPM3003E An unexpected database exception occurred on the snapshot database tables.](#)
- [HWNPM3004E The snapshot ID could not be found.](#)
- [HWNPM3500E The current transaction has been rolled back because of a deadlock.](#)
- [HWNPM3501E The current transaction has been rolled back because of a timeout.](#)
- [HWNPM3502E The current transaction has been rolled back because the database transaction log has been exhausted.](#)
- [HWNPM3503E The current transaction has been rolled back because the database disk space has been exhausted.](#)
- [HWNPM3600E The threshold identifier parameter value : threshold ID is not valid.](#)
- [HWNPM3601E The target component type parameter value : component type is not valid for the threshold identifier : threshold ID passed to the affected volumes and hosts reporting function.](#)
- [HWNPM3602E There was a problem retrieving the performance data needed to generate the affected volumes and hosts report for the device device name.](#)
- [HWNPM3603E The sample volume performance data needed to generate the affected volumes and hosts report for the device device name was not found in the IBM Spectrum Control database.](#)
- [HWNPM3604E There are no volumes associated with the specified target component, component name, in the IBM Spectrum Control database. Therefore, the resulting Affected Volumes and Hosts report will be empty.](#)
- [HWNPM4000E Unable to retrieve the device agent that managed this device: device identifier.](#)
- [HWNPM4001E Timeout while starting performance data collection for this device: device identifier.](#)
- [HWNPM4002E Unable to start performance data collection for this device: device identifier.](#)
- [HWNPM4003E Performance data collection has already been enabled for this device: device identifier.](#)
- [HWNPM4004E Failed to enable performance data collection for this device: device identifier.](#)
- [HWNPM4005I Successfully enabled performance data collection on the storage subsystem, using device access point SMI-S provider address.](#)
- [HWNPM4006E An exception occurred while starting performance data collection for this device: device identifier.](#)
- [HWNPM4007E A timeout occurred while stopping performance data collection for this device: device identifier.](#)
- [HWNPM4008E Unable to stop performance data collection for this device: device identifier.](#)
- [HWNPM4009E Performance data collection is not enabled for this device: device identifier.](#)
- [HWNPM4010E Failed to disable performance data collection for this device: device identifier.](#)
- [HWNPM4011I Successfully disabled performance data collection on the storage subsystem, using device access point SMI-S provider address.](#)
- [HWNPM4012E An exception occurred while stopping performance data collection for this device: device identifier.](#)
- [HWNPM4013E A timeout occurred while retrieving the status of the performance data collection for this device: device identifier.](#)
- [HWNPM4014E Unable to retrieve the status of the performance data collection for this device: device identifier.](#)
- [HWNPM4015I Performance data collection is not enabled for this device: device identifier.](#)
- [HWNPM4016I Performance data collection is enabled for this device: device identifier.](#)
- [HWNPM4017E Unable to determine the status of the performance data collection for this device: device identifier.](#)
- [HWNPM4018E Failed to retrieve the status of the performance data collection for this device: device identifier.](#)
- [HWNPM4019E A timeout occurred while polling the performance statistics for this device: device identifier.](#)
- [HWNPM4020E Unable to retrieve the performance statistics for this device: device identifier.](#)
- [HWNPM4021E No performance statistics available at the current time for this device: device identifier.](#)
- [HWNPM4022E Failed to disable performance data collection for this device: device identifier.](#)
- [HWNPM4023W A set of performance statistics data was empty for this device: device identifier.](#)
- [HWNPM4024E An exception occurred while stopping performance data collection for this device: device identifier.](#)
- [HWNPM4025E Unable to retrieve storage subsystem for this device: device identifier.](#)
- [HWNPM4026E Failed to retrieve storage subsystem for this device: device identifier.](#)
- [HWNPM4027E Failed to properly initialize counter data service for this device: device identifier.](#)

- [HWNPM4028W Performance data cannot be collected because the security role authority of the user account user name for accessing device identifier is not sufficient.](#)
- [HWNPM4029W Performance data cannot be collected because the collection of performance statistics is stopped on device identifier. The security role authority of the user account user name for accessing the storage system is not sufficient to start the collection of performance statistics.](#)
- [HWNPM4030W Performance data cannot be collected. The performance interval device interval on device identifier is greater than the sample interval and the set security role authority of the user account user name is not sufficient to update the interval value on device identifier .](#)
- [HWNPM4051E Failed to obtain a reference to the Performance Manager Configuration Data Service for this device: device name.](#)
- [HWNPM4052E Error occurred in trying to retrieve a device agent for this device: device name.](#)
- [HWNPM4053E Unable to locate or retrieve the device agent that manages this device: device name.](#)
- [HWNPM4054E Error occurred in trying to construct the poll state information for this device: device name.](#)
- [HWNPM4055E Unable to construct the poll state information for this device: device name.](#)
- [HWNPM4056E SMI-S provider operation triggered a timeout \(step timeout= step timeout value seconds, operation timeout= total timeout value seconds.\).](#)
- [HWNPM4057E Mismatch in device identifier for this device: device name.](#)
- [HWNPM4058E Failed to build the parameter Map for this device: device name.](#)
- [HWNPM4059I Performance data collection has already been enabled for this device: device name.](#)
- [HWNPM4060I Performance data collection was successfully started for this device: device name.](#)
- [HWNPM4061E Performance data collection could not be started for this device: device name.](#)
- [HWNPM4062I Performance data collection successfully stopped for this device: device name.](#)
- [HWNPM4063W Parse exception in performance data collected this device: device name.](#)
- [HWNPM4064E Wrong format in performance data collected for this device: device name.](#)
- [HWNPM4065W number of null time stamps null time stamp\(s\) for performance data collected from the device were substituted by server time stamp\(s\).](#)
- [HWNPM4066W count of null operational status null Port Operational Status value\(s\) for performance data collected from the device was/were substituted by default value\(s\).](#)
- [HWNPM4081E A database cursor operation failed.](#)
- [HWNPM4082E A database connect operation failed.](#)
- [HWNPM4083E A database retrieve operation failed.](#)
- [HWNPM4084E A database operation failed.](#)
- [HWNPM4085E A database query operation failed.](#)
- [HWNPM4086W A database query gave no result rows.](#)
- [HWNPM4087W Missing or invalid association between SMI-S provider SMI-S provider URL and device device name. The configured SMI-S provider is inoperative, or may no longer be managing the specified device.](#)
- [HWNPM4091E Encountered an error during execution of a discover service process.](#)
- [HWNPM4092E Encountered exception during execution of a discover service process.](#)
- [HWNPM4093E An input business object could not be converted to a CIMInstance.](#)
- [HWNPM4100E Failed to initialize SVC counter data service discover service reference.](#)
- [HWNPM4101E Failed to initialize SVC counter data service configuration service reference.](#)
- [HWNPM4102E Failed to parse performance data file time stamp suffix: filename.](#)
- [HWNPM4103E SMI-S provider operation timeout \(timeout value seconds\) expired.](#)
- [HWNPM4104E Failed to retrieve SMI-S provider password for SVC counter data service access point: access point.](#)
- [HWNPM4105E Encountered an error when communicating with the device agent.](#)
- [HWNPM4106E Encountered invalid SVC component type: component type.](#)
- [HWNPM4107E Failed to create performance data object: performance data object class.](#)
- [HWNPM4108E TimeZone property is not defined for SVC cluster: cluster identifier.](#)
- [HWNPM4109E SVC cluster TimeZone property is set to unrecognized value: timezone id and name.](#)
- [HWNPM4110E StatisticsStatus property is not defined for SVC cluster: cluster identifier.](#)
- [HWNPM4111E Failed to retrieve dump filename dump from SVC node node identifier \(return code = return code\).](#)
- [HWNPM4112E IsConfigNode property is not defined for SVC node: node identifier.](#)
- [HWNPM4113E Caught exception while processing SVC XML performance data.](#)
- [HWNPM4114E SVC cluster cluster identifier has more than one configuration node.](#)
- [HWNPM4115E SVC cluster cluster identifier does not have a configuration node.](#)
- [HWNPM4116W Failed to associate SVC performance data from non-configuration node with SVC performance data from configuration node.](#)
- [HWNPM4117W Encountered incomplete SVC performance data sample.](#)
- [HWNPM4118E Firmware version information is not available for storage subsystem subsystem name. Performance data collection cannot proceed.](#)
- [HWNPM4119E The firmware installed on storage subsystem subsystem name \(firmware version\) is not supported for performance data collection. The minimum level of firmware supported for performance data collection is firmware version.](#)
- [HWNPM4150E Unable to retrieve storage subsystem for this device: device identifier.](#)
- [HWNPM4151E Unable to determine the status of any performance data collection for this device: device identifier.](#)
- [HWNPM4152E Performance data collection has already been enabled for this device: device identifier.](#)
- [HWNPM4153E Performance data collection is not enabled for this device: device identifier.](#)
- [HWNPM4154E Unable to start performance data collection for this device: device identifier.](#)
- [HWNPM4155E Failed to enable performance data collection for this device: device identifier.](#)
- [HWNPM4156E Unable to stop performance data collection for this device: device identifier.](#)
- [HWNPM4157E Failed to disable performance data collection for this device: device identifier.](#)
- [HWNPM4158E Unable to complete start performance data collection task for this device: device identifier.](#)
- [HWNPM4159E Unable to complete stop performance data collection task for this device: device identifier.](#)
- [HWNPM4160E Unable to complete performance data collection status query task for this device: device identifier.](#)
- [HWNPM4161E Performance data collection is not enabled for this device: device identifier.](#)
- [HWNPM4162E Unable to retrieve port performance statistics data for this device: device identifier.](#)
- [HWNPM4163E Unable to retrieve volume performance statistics data for this device: device identifier.](#)
- [HWNPM4164E Unable to retrieve rank performance statistics data for this device: device identifier.](#)
- [HWNPM4165E Unable to retrieve performance statistics data for this device: device identifier.](#)
- [HWNPM4166E Unable to complete polling for performance data collection task for this device: device identifier.](#)
- [HWNPM4167E Unable to retrieve a device agent for this device: device identifier.](#)
- [HWNPM4168E Failed attempt to use device device identifier counter data service with device different device identifier.](#)
- [HWNPM4169E An invalid access point of device agent URL was used to acquire the agent for this device: device identifier.](#)
- [HWNPM4170E The device agent's configuration for device identifier has changed from the given access point, device agent URL.](#)
- [HWNPM4171E Performance data collection start task timed out after time seconds for device: device identifier.](#)
- [HWNPM4172E Performance data collection stop task timed out after time seconds for device: device identifier.](#)
- [HWNPM4173E Performance data collection check status task timed out after time seconds for device: device identifier.](#)
- [HWNPM4174E Performance data collection poll task timed out after time seconds for device: device identifier.](#)

- [HWNPM4175W An error occurred while parsing statistics for port port identifier. Its statistics will be excluded.](#)
- [HWNPM4176W An error occurred while parsing statistics for volume volume identifier. Its statistics will be excluded.](#)
- [HWNPM4177W An error occurred while parsing statistics for rank rank identifier. Its statistics will be excluded.](#)
- [HWNPM4178E Failed to decrypt the device agent's password for device device identifier.](#)
- [HWNPM4179W Performance data collection is currently enabled with errors for device device identifier.](#)
- [HWNPM4180E Unable to retrieve key identifier value from the internal discover process.](#)
- [HWNPM4181W number of ports of the port statistics from the device agent were unrecognized and were not included in this sample interval.](#)
- [HWNPM4182W number of volumes of the volume statistics from the device agent were unrecognized and were not included in this sample interval.](#)
- [HWNPM4183W number of ranks of the rank statistics from the device agent were unrecognized and were not included in this sample interval.](#)
- [HWNPM4184E The device agent configured for this storage subsystem is not supported for this task. The current version, version number, is downlevel from the minimum required, version number.](#)
- [HWNPM4185W The device agent did not return all performance statistics data for this time interval. The incomplete data is being processed.](#)
- [HWNPM4186W The ESS SMI-S provider did not return performance statistics data for both clusters for this time interval. The incomplete data is being processed.](#)
- [HWNPM4187W The device does not support performance management for pool pool ID because it contains Space Efficient Volumes. Only incomplete performance data can be collected for array array ID.](#)
- [HWNPM4188W The performance monitor was unable to collect performance statistics data from the device agent for the following component types: component list.](#)
- [HWNPM4189W number of MDisks of the MDisk statistics from the device agent were unrecognized and were not included in this data collection interval.](#)
- [HWNPM4190W number of nodes of the node statistics from the resource agent were unrecognized and were not included in this data collection interval.](#)
- [HWNPM4191W number of modules out of total number of modules module statistics could not be retrieved from the device agent due to errors, and were not included in this data collection interval.](#)
- [HWNPM4192W number of Drives of the drive statistics from the device agent were unrecognized and were not included in this data collection interval.](#)
- [HWNPM4193W number of Volume-copies of the volume-copy statistics from the device agent were unrecognized and were not included in this data collection interval.](#)
- [HWNPM4194W number of partitions of the partition statistics from the device agent were unrecognized and were not included in this data collection interval.](#)
- [HWNPM4195W number of file systems of the file system statistics from the device agent were unrecognized and were not included in this data collection interval.](#)
- [HWNPM4250E Failed to start the discover service for the SMI-S counter data service.](#)
- [HWNPM4251E Failed to start the configuration service for the SMI-S counter data service.](#)
- [HWNPM4252I Successfully returned access point device name for device device name.](#)
- [HWNPM4253I Successfully stopped SMI-S counter data service on access point access point for device device name.](#)
- [HWNPM4254I The SMI-S counter data service is active on access point access point for device device name.](#)
- [HWNPM4255I The SMI-S counter data service is inactive on access point access point for device device name.](#)
- [HWNPM4256I Performance statistics successfully returned on access point access point for device device name.](#)
- [HWNPM4257W Performance statistics not returned on access point access point for device device name.](#)
- [HWNPM4258E No SMI-S providers found for device device name.](#)
- [HWNPM4259E No storage subsystem found for device device name.](#)
- [HWNPM4260E Failed to initialize the polling context for device device name.](#)
- [HWNPM4261E Failed to retrieve the device capabilities for device device name.](#)
- [HWNPM4262E A database exception occurred trying to retrieve the device capabilities for device device name.](#)
- [HWNPM4263E A database exception occurred trying to retrieve the storage subsystem for device device name.](#)
- [HWNPM4264W Failed to retrieve manifest for component type.](#)
- [HWNPM4265E A database exception occurred trying to retrieve the Manifests for device device name.](#)
- [HWNPM4266E No manifests found for device device name.](#)
- [HWNPM4267E A database exception occurred trying to retrieve the discovery parameters for device device name.](#)
- [HWNPM4268E Statistics record not correctly formatted due to exception local exception string.](#)
- [HWNPM4269E Statistics record not correctly parsed due to exception local exception string.](#)
- [HWNPM4270W The block storage statistics is not formatted for device device name.](#)
- [HWNPM4271E The SMI-S provider found for device device name is not valid.](#)
- [HWNPM4272E The storage subsystem found for device device name is not valid.](#)
- [HWNPM4273W Discarding the stale performance statistics returned on access point access point for device device name.](#)
- [HWNPM4274E The SMI-S provider found for this device has changed. Please re-run SMI-S provider discovery and probe.](#)
- [HWNPM4300E Access to the agent or device has been denied. Ensure that valid credentials have been specified for agent agent name.](#)
- [HWNPM4301E The device or device agent did not respond within the allotted time \(timeout value seconds\).](#)
- [HWNPM4302E New performance data is not yet available for the device. Statistics with time stamps later than time stamp could not be found.](#)
- [HWNPM4303E An agent API call \(API name\) failed while attempting to retrieve performance data for the device.](#)
- [HWNPM4304E The request for performance data could not be retrieved from the queue by the data collector.](#)
- [HWNPM4305W No samples were received from the data collector in the expected time. The data might still arrive automatically after connection is recovered.](#)
- [HWNPM4306E The data collector failed to connect to the storage management service because of invalid credentials. No performance manager data can be collected from device name until valid credentials are available.](#)
- [HWNPM4502E Attempt to delete a default policy.](#)
- [HWNPM4503E A database update operation failed.](#)
- [HWNPM4504E A database insert operation failed.](#)
- [HWNPM4505E A database delete operation failed.](#)
- [HWNPM4506E A database cursor operation failed.](#)
- [HWNPM4507E A database connect operation failed.](#)
- [HWNPM4508E A database retrieve operation failed.](#)
- [HWNPM4509E A database operation failed.](#)
- [HWNPM4510E A database query operation failed.](#)
- [HWNPM4511E A database commit operation failed.](#)
- [HWNPM5200E The performance manager failed to publish event even name due to exception exception.](#)
- [HWNPM5210E The performance manager failed to receive event from other modules.](#)
- [HWNPM5211E The first parameter passed to this method is null.](#)
- [HWNPM5212E The second parameter passed to this method is invalid.](#)
- [HWNPM5400E The performance data collection identifiers are not valid integers: schedule ID {0}, schedule run number {1}, job run number {2}.](#)
- [HWNPM5401E There was a problem establishing the database connection: {0}.](#)
- [HWNPM5402E There was a problem creating the new run job entry: {0}.](#)

- [HWNPM5403E There was a problem updating the run job entry {0}; {1}.](#)
- [HWNPM5404E There was a problem closing the database connection: {0}.](#)
- [HWNPM5405E There was a problem inserting a new run job into the database: {0}.](#)
- [HWNPM5406E There was a problem executing an update for run job number {0} in the database.](#)
- [HWNPM5407E There was a problem executing an update for run job number {0} in the database.](#)
- [HWNPM5408E There was a problem executing an update for run number {0} in the database.](#)
- [HWNPM5409I Successfully retrieved the configuration data for the elastic device. Found number of nodes Nodes and number of file systems File systems.](#)
- [HWNPM5410W The performance monitor could not collect performance data for the following cluster nodes: nodes names.](#)
- [HWNPM5411W The performance monitor could not collect performance data for the following filesystems: filesystem names.](#)
- [HWNPM5412E Performance statistics collection is not enabled.](#)
- [HWNPM5413E The process failed because the userid or password provided failed to connect to the Export Tool.](#)
- [HWNPM5414E The process failed because the Hitachi SVP was busy and did not return data or timed out.](#)
- [HWNPM5415E The process failed because the performance interval is set to 0.](#)
- [HWNPM5416E The process failed because the performance interval for the storage system is not supported.](#)
- [HWNPM5417E The process failed because the Hitachi VSP Model being monitored is not known. A Hitachi Export Tool to match it cannot be found.](#)
- [HWNPM5418E The process failed because the data collected is out of range.](#)
- [HWNPM5419E Performance data can't be collected.](#)

HWNPM0001E The specified summarization level (*level*) is invalid. It must be an integer value between *minimum* and *maximum*, inclusive.

Explanation

The summarization level that was passed as argument on the Performance Manager API call is invalid. The method that was called returns this error indicator.

Action

Modify the caller of this method to pass a valid summarization level as parameter. The valid levels are enumerated via the `summtpe_*` constants in the `com.ibm.tpc.perf.api.ApiConstants` interface.

HWNPM0002E The specified device category (*category*) is invalid. It must be an integer value between *minimum* and *maximum*, inclusive.

Explanation

The device category that was passed as argument on the Performance Manager API call is invalid. The method that was called returns this error indicator.

Action

Modify the caller of this method to pass a valid device category as parameter. The valid categories are enumerated via the `devcat_*` constants in the `com.ibm.tpc.perf.api.ApiConstants` interface.

HWNPM0003E The specified device type (*type*) is invalid. It must be an integer value between *minimum* and *maximum*, inclusive.

Explanation

The device type that was passed as argument on the Performance Manager API call is invalid. The method that was called returns this error indicator.

Action

Modify the caller of this method to pass a valid device type as parameter. The valid types are enumerated via the `devtype_*` constants in the `com.ibm.tpc.perf.api.ApiConstants` interface.

HWNPM0004E The specified component type (*type*) is invalid. It must be an integer value between *minimum* and *maximum*, inclusive.

Explanation

The component type that was passed as argument on the Performance Manager API call is invalid. The method that was called returns this error indicator.

Action

Modify the caller of this method to pass a valid component type as parameter. The valid types are enumerated via the `comtype_*` constants in the `com.ibm.tpc.perf.api.ApiConstants` interface.

HWNPM0006E The string specified as parameter (*string*) exceeded its allowed length (*maximum length*).

Explanation

The string parameter that was passed as argument on the Performance Manager API call is too long. The method that was called returns this error indicator.

Action

Length limitations usually arise due to the need for saving the given string into a fixed length database column. Modify the caller of this method to pass a shorter string as parameter.

HWNPM0007E The value specified as parameter (*value*) is invalid.

Explanation

The parameter value that was passed as argument on the Performance Manager API call is invalid. The method that was called returns this error indicator.

Action

Refer to the method or class documentation to determine the allowed values for the particular parameter. Modify the caller of this method to pass a valid value to the API.

HWNPM0008E A required parameter is missing (*null*).

Explanation

The parameter value that was passed as argument on the Performance Manager API call is NULL, but the parameter is required. The method that was called returns this error indicator.

Action

Refer to the method or class documentation to determine the allowed values for the particular parameter. Modify the caller of this method to pass a valid value to the API.

HWNPM0010E The specified device ID (*device ID*) is invalid. It must conform to the pattern '*name+nameFormat*'.

Explanation

The device identifier string that was passed as argument on the Performance Manager API call is invalid. The method that was called returns this error indicator.

Action

Modify the caller of this method to pass a valid device ID as parameter.

HWNPM0011E The specified component ID (*component ID*) is invalid. It must be a simple WWN (16 hexadecimal characters).

Explanation

The component identifier string that was passed as argument on the Performance Manager API call is invalid. The method that was called returns this error indicator.

Action

Modify the caller of this method to pass a valid component ID as parameter.

HWNPM0012E The specified component ID (*component ID*) was not found or is not unique in the IBM Spectrum Control database.

Explanation

The component identifier string that was passed as argument on the Performance Manager API call does not correspond to a device known to IBM Spectrum Control. The method that was called returns this error indicator.

Action

Modify the caller of this method to pass a valid component ID as parameter. The parameter string must uniquely identify the required component.

HWNPM0013E The specified component ID (*component ID*) is invalid.

Explanation

The component identifier string that was passed as argument on the Performance Manager API call is invalid. It must conform to the pattern 'compName+deviceName+deviceNameFormat'. The method that was called returns this error indicator.

Action

Modify the caller of this method to pass a valid component ID as parameter.

HWNPM0015E Failed to retrieve the requested data because the service is unavailable.

Explanation

The Device server is unable to connect to the database. The API call cannot be completed.

Action

To resolve the problem, check the following and then try the action again:

- Verify that the database manager is operational.
- Stop and start the product servers.

HWNPM0021E The device identifier specified as parameter (*device ID*) is invalid.

Explanation

The device identifier that was passed as an argument for the product API is invalid. The device identifier was either not found or is not unique in the database.

Action

Modify the caller of this method to pass a valid device identifier as a parameter. The parameter must uniquely identify the required device.

HWNPM0090E Failed to retrieve the requested data because the service is unavailable.

Explanation

The Device server was unable to complete a database operation. This might have caused the requested operation to fail.

Action

To resolve the problem, check the following and then try the action again:

- Verify that the database manager is operational.
- Stop and start the product servers.

HWNPM0099E The requested operation failed because of an internal error.

Explanation

An internal error occurred during the requested operation. Internal errors can be caused by a database inconsistency or corruption, or can be caused by a programming error.

Action

Try the operation again. Check for error messages in the log files for the Device server. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM0101E Unable to create the specified performance service instance ({0}).

Explanation

The Performance Manager failed trying to create a service instance (performance collection, performance threshold, or performance reporting service). The accompanying exception object should give more clues as to the exact problem.

Action

If you have made changes to the PM configuration file, reverse those changes. Try to restart the IBM Spectrum Control Device Server service.

If you are overriding the default instantiation class by specifying the associated property as argument to the factory class, ensure that the specified instantiation class exists and is accessible by the JVM.

HWNPM0200I This operation (*operation name*) on Performance Manager was successful.

Explanation

None.

Action

None.

HWNPM0201E The device (*device_id*) that was passed to the method is invalid.

Explanation

The device is null or is not supported.

Action

Provide the correct device to the method.

HWNPM0202E The device category (*device_category*) that was passed to the method is invalid.

Explanation

The device category is null or not supported.

Action

Provide the correct device category to the method.

HWNPM0203E The device type received (*device_type*) is invalid.

Explanation

The device type is null or not supported.

Action

Provide the correct device type to the method.

HWNPM0204E The device type - HOST - that was passed to the method is not supported.

Explanation

The HOST device type is not supported for performance data collection.

Action

None.

HWNPM0205E The specified performance collection policy is invalid.

Explanation

The specified policy is invalid and could not be successfully set into effect. The previous performance collection policy remains in effect for the associated device(s).

Action

Please ensure that all performance collection rules were followed when the new policy was created.

HWNPM0209I The device type and device category are valid.

Explanation

None.

Action

None.

HWNPM0210E Collector failed to start due to system failure.

Explanation

A system failure caused a collector not to start.

Action

None.

HWNPM0220E Collector failed to stop due to system failure.

Explanation

A system failure caused a running collector to fail.

Action

None.

HWNPM0230E One or more of the specified performance collection policies are invalid.

Explanation

A performance collection policy object that was specified as an argument for the product API call is invalid. The called method could not process the policy object.

Action

Modify the caller of this method to pass valid performance collection policy objects as arguments. One common cause of problems is the specification of an invalid device identifier for custom policies. Ensure that the policy device ID field corresponds to a device that is valid and known to the product.

HWNPM0231W The specified performance collection policy is ignored because it conflicts with another policy in the same parameter list.

Explanation

The performance collection policy object that was specified as argument on the Performance Manager API call is ignored, because another policy specified in the same call is effectively equivalent to this policy. Equivalent policies have the same device type and device ID. The other policy has been used for the purpose of this call, and this policy has been ignored.

Action

Modify the caller of this method to remove any conflicting or duplicate policy objects.

HWNPM0232E The specified performance collection policy contains an unsupported interval length.

Explanation

The performance collection policy object that was specified as argument on the Performance Manager API call is invalid because it contains an unsupported interval length attribute. The supported interval lengths can vary for individual devices, and for individual device types.

Action

Consult the performance collection rules to determine the supported interval lengths for devices of the desired type, and/or the particular device in question. Then modify the caller of this method to specify a valid interval length attribute in the passed policy object.

HWNPM0233E The specified performance collection policy contains an unsupported frequency.

Explanation

The performance collection policy object that was specified as argument on the Performance Manager API call is invalid because it contains an unsupported frequency attribute. The supported frequencies can vary for individual devices, and for individual device types.

Action

Consult the performance collection rules to determine the supported frequencies for devices of the desired type, and/or the particular device in question. Then modify the caller of this method to specify a valid frequency attribute in the passed policy object.

HWNPM0234E The specified performance collection policy contains an unsupported duration.

Explanation

The performance collection policy object that was specified as argument on the Performance Manager API call is invalid because it contains an unsupported duration attribute. The supported durations can vary for individual devices, and for individual device types.

Action

Consult the performance collection rules to determine the supported durations for devices of the desired type, and/or the particular device in question. Then modify the caller of this method to specify a valid duration attribute in the passed policy object.

HWNPM0240E The attempt to update the specified performance collection policies has failed.

Explanation

An internal error occurred while performing the requested IBM Spectrum Control Performance Manager API operation. Internal errors can be caused by a database inconsistency or corruption, or can be due to a programming error.

Action

Please retry the failing operation. If the failure persists, please contact IBM software support. More details about the exact failure will be available in the Performance Manager trace logs in the device server log directory.

Related reference

- [Getting support](#)

HWNPM0241E The attempt to reset the specified performance collection policies has failed.

Explanation

An internal error occurred while performing the requested IBM Spectrum Control Performance Manager API operation. Internal errors can be caused by a database inconsistency or corruption, or can be due to a programming error.

Action

Please retry the failing operation. If the failure persists, please contact IBM software support. More details about the exact failure will be available in the Performance Manager trace logs in the device server log directory.

Related reference

- [Getting support](#)

HWNPM0242E The attempt to remove the specified performance collection policies has failed.

Explanation

An internal error occurred while performing the requested IBM Spectrum Control Performance Manager API operation. Internal errors can be caused by a database inconsistency or corruption, or can be due to a programming error.

Action

Please retry the failing operation. If the failure persists, please contact IBM software support. More details about the exact failure will be available in the Performance Manager trace logs in the device server log directory.

Related reference

- [Getting support](#)

HWNPM0249W An attempt to dynamically update one or more running performance collectors with a new performance collection policy has failed.

Explanation

When a new performance collection policy is set into effect, any running performance collectors using that policy must be dynamically updated, for the new settings to take effect. This dynamic update failed for one or more of the running collectors associated with a particular policy.

Action

The policy has been successfully saved in the database, so the operation does not need to be repeated. However the running performance collector needs to be manually stopped and restarted for the new settings to take effect.

HWNPM0250E One or more default performance collection policies are missing from the database.

Explanation

One default performance collection policy must exist in the database for each supported device type. The IBM Spectrum Control Performance Manager will not function properly without these default policies. The fact that they are missing from the database indicates a database corruption of some kind.

Action

Please retry the failing operation. If the failure persists, please contact IBM software support. If a database corruption seems plausible, and you have a recent backup of the IBM Spectrum Control database, you can also try shutting down IBM Spectrum Control and restoring the old database. However remember that you will lose all information that has been added to the database since the backup was made.

Related reference

- [Getting support](#)

HWNPM0281I Performance monitoring is unavailable for resource *resource_name* because an agent for monitoring the resource was not defined. For IBM Spectrum Scale, the problem might occur because the data collection service cannot connect to port 9084 on the node where the collector component of the IBM Spectrum Scale performance monitoring tool is running.

Explanation

For the product to collect performance statistics and other necessary information to monitor a resource, an agent that manages the resource must be defined. An agent can be an SMI-S provider, ZIMon agent for IBM Spectrum Scale, or a native API for a resource.

Action

Define a data source for the resource, such as an SMI-S provider, a ZIMon agent for IBM Spectrum Scale, or a resource-specific service.

Before you can collect performance data for IBM Spectrum Scale, the IBM Spectrum Control server must be able to connect to port 9084 on the node where the collector component of the IBM Spectrum Scale performance monitoring tool is running.

To ensure that the IBM Spectrum Control server can connect to the collector node, complete the following steps:

1. Determine the node that is configured as the collector node by viewing the `/opt/IBM/zimon/ZIMonSensors.cfg` file on one of the sensor nodes. The collector node is set in the host property of the collectors section in this file. For example, `collectors = {host = "node3" port = "4739"}`
2. Ensure that the host property is set to one of the following options:
 - An IP address that can be reached by the IBM Spectrum Control server
 - A host name that can be resolved to a reachable IP address by the IBM Spectrum Control server

Related reference

- [🔗 Configuring the collection of performance data for IBM Spectrum Scale](#)

HWNPM0282I Performance monitoring is unavailable for resource *resource_name* because the associated data sources are unable to collect performance data from the resource.

Explanation

One or more data sources for the resource are defined, but the particular types of data sources cannot collect performance data. The data sources might be able to complete other tasks for the resource.

Action

Add the resource using a data source that can retrieve performance data from the resource, if such a data source is available.

HWNPM0283I Performance monitoring is unavailable for *resource_name* because this resource does not support performance monitoring.

Explanation

One or more agents for the resource are defined in the product, but these agents cannot retrieve performance statistics for the resource. The agents might be able to perform other tasks for the resource.

Action

Upgrade or reinstall the firmware for the device.

HWNPM0284I Performance monitoring is unavailable for resource *resource_name* because the associated agent does not have the required level of software *agent_level*.

Explanation

The product does not support performance monitoring of the device at its current agent software level. In most cases, the current agent level is earlier than the required level, in which case an upgrade of the agent software is necessary. In rare cases, the agent software level is a later level that is not supported.

Action

Upgrade or reinstall the agent software for the device.

HWNP0285I Performance monitoring is unavailable for resource *resource name* because the associated agent is unable to fully monitor the resource.

Explanation

A device agent (an SMI-S provider, SMI agent, or a proprietary service API, etc.) must exist for IBM Spectrum Control to be able to retrieve the performance statistics and other necessary information to monitor the device. The currently defined agent(s) for the device will not allow IBM Spectrum Control to perform performance monitoring on the device. This can happen primarily for FC switches in mixed fabrics if there is a mismatch between the SMI agent and device, in other words if one switch vendor's SMI agent was able to discover a different vendor's switch. The mismatching agent will not be able to fully manage the other vendor's switches, and in particular, will not be able to retrieve performance statistics for those switches.

Action

Please define a "matching" agent for the particular vendor's device, for example for Brocade switches ensure that there is a Brocade SMI agent defined to manage the switch, for McData switches a McData SMI agent, and so forth. Switches with embedded SMI agents can only be managed using that embedded agent.

HWNP0286I Performance monitoring is unavailable for resource *resource name* because the associated SMI-S provider does not have the required SMI-S support.

Explanation

An SMI-S provider has been defined in IBM Spectrum Control to manage the device, but unfortunately that agent does not have the necessary SMI-S support to allow retrieval of performance statistics for the device. For FC switches, the SMI-S provider needs to support the SMI-S "Switch" profile, and must support the SMI-S level 1.0.2 at minimum. For storage subsystems, the SMI-S provider needs to support the "Array" profile and the "Block Server Performance" (BSP) sub-profile.

Action

Install an SMI-S provider with the necessary SMI-S support, and configure it to manage the device in question. Define the agent to IBM Spectrum Control using the IBM Spectrum Control GUI Administrative Services, Data Sources. Note that some vendors require purchase of a separate feature to enable this support.

HWNP0287I Performance monitoring is unavailable for resource *resource name* because the resource or the associated agent does not support performance monitoring.

Explanation

This can happen primarily for virtual devices, for example a logical FC switch. IBM Spectrum Control does not support performance data collection for virtual devices.

Action

Run the performance monitor on the physical device rather than on the individual virtual devices configured for the physical device.

HWNP0288I Performance monitoring is unavailable for resource *resource name* because the resource was not probed.

Explanation

To monitor the performance of a resource, you must first run a probe to collect asset and configuration information about that resource. Probes are data collection jobs that you schedule or start immediately.

Action

Define and run a probe for the resource. To run a probe, go the resource list page for the resource, right-click the resource, and select Data Collection > Start Probe. For example, to probe a storage system, go to Storage Systems page, right-click the storage system, and select Data Collection > Start Probe. After the probe completes successfully, run the performance monitor again.

To monitor the performance of a switch, the following conditions must be met:

- The switch must be probed through a SMI or SNMP agent. Probing a switch through a Storage Resource agent does not provide sufficient performance information.
- The probe of a SMI agent must complete successfully and discover at least one port on the switch. If the probe does not discover one or more ports, the switch is not eligible for performance monitoring using that SMI agent. If multiple agents are used for a probe, it is possible that an SNMP or Storage Resource agent discovers switch ports but a SMI agent does not. If you suspect that this situation occurred, stop and restart the SMI agent and rerun the fabric probe.

HWNPM0289W Performance monitoring is unavailable for resource *resource_name* because no agents are available.

Explanation

One or more of the agents that monitor the resource are not available. If multiple agents are monitoring the resource, the product tries to use each agent to collect performance data. If none of the agents are up and running, no data can be collected.

Action

To resolve the problem, try the following actions:

- Ensure that the agent for the resource is up and running.
- Verify that the local area network is available.
- Verify that the product can communicate with the agent.
- Restart the performance monitor that collects data about the resource.
- Check for error messages in the log files for the Device server.

HWNPM0290E Performance monitoring is unavailable for resource *resource_name* because the associated agent was could not be selected.

Explanation

This is due to an internal failure in the device server.

Action

Try to delete and then recreate the performance monitor job definition. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM0291I Performance monitoring is unavailable for switch *resource_name* because the switch has no ports.

Explanation

All performance data for FC switches is associated with the individual switch ports. If the switch has no ports, no performance data can be collected for it.

Action

If the switch does have ports but the ports were not discovered, run the probe for the switch again. To run a probe, go to the resource list page for switches, right-click the switch, and select Data Collection > Start Probe. After the probe completes successfully, run the performance monitor again.

HWNPM0292I Performance monitoring is unavailable for switch *resource_name* because the switch was not probed using the correct agent.

Explanation

To monitor the performance of a resource, you must first run a probe to collect asset and configuration information about that resource. Probes are data collection jobs that you schedule or start immediately. In addition, to successfully collect performance data for FC switches using SMI agents, it is necessary to have probed the switch

using that exact SMI agent. If the switch was probed using a Storage Resource agent or SNMP agent, it is not possible to collect performance data using a SMI agent. If the switch is managed by multiple SMI agents, only those SMI agents that have been previously used to probe the switch can be used to collect performance data for it.

Action

If a probe is not yet defined for the switch, define and run a probe for the resource. To run a probe, go to the resource list page for switches, right-click the switch, and select Data Collection > Start Probe. After the probe completes successfully, run the performance monitor again. If the Data Collection -> Start Probe action is not available, a probe is not yet defined. Right-click the switch and select Data Collection -> Schedule to define and save the probe parameters, and then select Data collection -> Start Probe to run it.

Even if a probe has successfully completed in the past, it may be necessary to run the probe again. When new agents have been added for managing the switch, or old agents have become unreachable, running a new probe will ensure that all the data needed for successful performance monitoring was collected.

If running a new probe does not resolve the error, it may be necessary to reset the SMI agent. It may not be obvious when a SMI agent is returning incorrect or incomplete data, especially when using multiple, different types of agents to probe the switch. Stop and restart the SMI agent and run the probe again.

HWNPM0293I Performance monitoring is unavailable for FlashSystem resource_name because its SNMP agent is disabled. You can enable SNMP for FlashSystem storage systems in the FlashSystem GUI<a>.

Explanation

Collection of performance data from a FlashSystem 840 or FlashSystem 900 requires the use of an SNMP data source. If the FlashSystem SNMP agent is not enabled, the performance monitor for that system fails with this message.

Action

Go to the management GUI of the FlashSystem storage system, then go to the SNMP Settings page, set the SNMP v1 Read Community name, and enable the SNMP agent. The SNMP agent for the FlashSystem storage system supports only SNMPv1 for authentication. If using SNMPv1 for authentication is unacceptable in your environment due to security reasons, do not enable the SNMP agent. If the SNMP agent is not enabled you cannot collect performance data from your FlashSystem storage system.

If IBM Spectrum Control is already monitoring the FlashSystem storage system, run another probe after you enable the SNMP Agent. To run a probe, go the Storage Systems page, right-click the storage system, and click Data Collection > Start Probe. After the probe completes successfully, run the performance monitor again.

HWNPM0300E There is an exception for each device processed in a multiple devices call.

Explanation

Because each device got an exception during the process, the method call will throw the first exception in the result array as the general exception for multiple devices.

Action

None.

HWNPM0390E A system failure occurred.

Explanation

A system failure occurred, which caused the operation to fail.

Action

None.

HWNPM0400I This operation (operation name) on Threshold Service was successful.

Explanation

None.

Action

None.

HWNPM0401E The device (*device_id*) that was passed to the method is invalid.

Explanation

The device is null or is not supported.

Action

Provide the correct device to the method.

HWNPM0410E The Performance threshold policy that was passed to the method (*threshold policy*) is null.

Explanation

The specified policy is invalid and could not be successfully set into effect. The previous performance threshold policy remains in effect for the associated device(s).

Action

Please ensure that all applicable rules were followed when the new policy was created.

HWNPM0411E The Performance threshold that was passed to the method (*threshold*) is null.

Explanation

The performance threshold is null.

Action

Provide a valid performance threshold to the method.

HWNPM0412E The Performance threshold filter that was passed to the method (*filter*) is null.

Explanation

The performance threshold filter is null.

Action

Provide a valid performance threshold filter to the method.

HWNPM0420E The device type received (*device_type*) is invalid.

Explanation

A device type is null or not supported by the IBM Spectrum Control system.

Action

Provide the correct device type to the method.

HWNPM0421E There is no default performance threshold policy or default threshold filter for this device.

Explanation

A default performance threshold policy and threshold filter should exist. These values are shipped with the product.

Action

Re-install the database.

HWNPM0425E There is an exception for each device processed in a multiple devices threshold call.

Explanation

Because each device received an exception during the process, the method call will throw the first exception in the result array as the general exception for multiple devices.

Action

None.

HWNPM0590E Performance Manager failed due to system failure.

Explanation

A system failure occurred which caused the failure of the operation.

Action

None.

HWNPM0600E Parameter number *a number* of the call made to the IBM Spectrum Control Performance Manager reporting API *method name of the api* is invalid. The invalid value is the invalid value of parameter *descriptive name of the parameter*.

Explanation

The caller of a Performance Manager reporting API receives a TPCMessage object containing this message if an API call fails with an InvalidParameterException. The message states what the invalid parameter is.

Action

The InvalidParameterException indicates the caller of the API passed an invalid value in one of the parameters. Usually this indicates an error in the implementation of the calling function. You will need to report the failure to the personnel supporting that function.

HWNPM0601E A request to continue the data retrieval can not be performed. Information for continuing the data retrieval does not exist.

Explanation

The caller of a Performance Manager reporting API receives a TPCMessage object containing this message, as part of an InvalidParameterException. The message states what the invalid parameter is.

Action

The exception might have occurred because the continuation state information has been deleted. The information might have been deleted by an explicit close request, or because a time limit exceeded because the previously initiated request, or because the continuation parameter was set incorrectly by the application.

HWNP0602E Support for the device type *device type name* is not available in the Performance Manager reporting API function *method name*.

Explanation

A caller invoked the specified Performance Manager reporting API function. Device-type specific support needed by that function is not currently available in IBM Spectrum Control.

Action

Report the problem to IBM support.

Related reference

- [Getting support](#)

HWNP0603E The performance reporting API *method_name* failed at *time_of_failure* as a result of an internal processing exception. The Performance Manager logs contain message *message_ID* that describes the internal processing exception.

Explanation

The caller of the reporting API receives a TPCMessage object containing this message if an API call fails with a TPCTServerException.

Action

The TPCTServerException indicates a condition beyond the control of the caller of the API. In some cases, retrying the API call might succeed. In other cases, you must report the failure to the personnel supporting your product. Include the information contained in this message in your report.

HWNP0604E The sort order parameter of a call to the performance reporting API *method_name* contains a value not included in the report columns list, at position *list_item* in the sort order list.

Explanation

The caller of a product reporting API receives a TPCMessage object containing this message if an API call fails with an InvalidParameterException. A value in the sort order list does not appear in the report columns list as is required.

Action

The metrics identified to be sorted on must be included in the list of report column metrics.

HWNP0605E Performance metadata cannot be displayed.

Explanation

This problem might occur when metadata is not available for the selected time range or device, or when the cache, cookies, and history in your web browser need to be cleared to refresh the GUI with the latest metadata.

Action

In some cases, you can clear the cache, cookies, and history in your web browser to resolve the problem.

In Mozilla Firefox:

1. Click Tools > Options and select Privacy and Security.
2. Under Cookies and Site Data, click Clear Data.
3. Under History, click Clear History.

In Google Chrome:

1. Click Settings and select Privacy and security.
2. Click Clear browsing data, accept the default choices, and click Clear data.

After the cache, cookies, and history is cleared, click the browser refresh button to update the product GUI.

For more detailed and up-to-date instructions about how to clear the cache, cookies, and history in different types and versions of web browsers, refer to the documentation for those browsers.

HWNPM0606E Unable to instantiate performance reporting service service class name.

Explanation

A failure occurred trying to instantiate the performance reporting service of Performance Manager. This usually indicates an internal error.

Action

If you have made changes to the PM configuration file, reverse those changes. Try to restart the IBM Spectrum Control Device Server service. If the problem persists, contact your IBM support representative.

Related reference

-  [Getting support](#)

HWNPM0607E An error occurred while the performance data was being retrieved.

Explanation

An error occurred that caused a delay in presenting the performance data that was being retrieved for the storage system and its resources.

Action

Wait a few minutes and try again. If you still can't see the performance data, go to the IBM Support page (<https://www.ibm.com/mysupport/>) where you can chat with an expert, browse troubleshooting topics and forums, open support cases, and access IBM Documentation.

Related reference

-  [Getting support](#)

HWNPM0630E An invalid operator (*operator identifier*) was specified for the filter expression.

Explanation

The caller of the performance reporting filter related constructor or method specified an invalid operator identifier.

Action

The application should use a valid operator identifier as described in the documentation of the method or constructor.

HWNPM0631E An invalid first operand (*operand class*) was specified for the filter expression. It must be a *operand class* class.

Explanation

The caller of the performance reporting filter related constructor or method specified an invalid first operand. The operand should be of the indicated type.

Action

The application should use a valid operand type as described in the documentation of the method or constructor.

HWNPM0632E An invalid first operand was specified for the filter expression. The data type of the operand (*data type*) is invalid or unsupported.

Explanation

The caller of the performance reporting filter related constructor or method specified an invalid first operand. The operand, usually an ICounter, should have a valid data type.

Action

The application should use a valid data type as described in the documentation of the method or constructor.

HWNPM0633E An invalid second operand (*operand class*) was specified for the filter expression. It must be a *operand class* class.

Explanation

The caller of the performance reporting filter related constructor or method specified an invalid second operand. The operand should be of the indicated type.

Action

The application should use a valid operand type as described in the documentation of the method or constructor.

HWNPM0650E The IBM Spectrum Control Performance Manager reporting API *method name* failed as a result of exception *the related exception* from a call to method *method name*, of the *lower-level service name* internal service.

Explanation

The Performance Manager reporting API received an exception from a call to an internal processing function. The exception resulted in failure of the reporting API call.

Action

Examine the log entries from the specified internal service for aid in determining the cause of the received exception.

HWNPM0651E The configuration data needed to generate the affected volumes and hosts report for the device *device name* was not found in the IBM Spectrum Control database.

Explanation

The configuration data of the device at the time the threshold or constraint violation occurred is not present in the database. The affected volumes and hosts report cannot be generated in absence of the device configuration data.

Action

Contact your IBM support representative.

Related reference

-  [Getting support](#)

HWNPM0652E The requested performance metrics are no longer available. Clear your browser cache to proceed.

Explanation

The performance metrics are not available.

Action

Clear your web browser cache and cookies. To learn how, see your web browser documentation .

HWNPM2000I Performance monitoring is enabled.

Explanation

The ability to monitor the performance of resources is enabled.

Action

To monitor the performance of a storage system or switch, schedule a performance monitor for that resource. Performance monitors can collect performance data daily and enable you to track trends and identify performance bottlenecks.

HWNPM2001E The IBM Spectrum Control Performance Manager is not operational.

Explanation

The Performance Manager component of IBM Spectrum Control is not running or has been stopped.

Action

Restart the Performance Manager Service and try again.

HWNPM2002E An initialization error occurred.

Explanation

A failure occurred during the initialization of a service of the Device server.

Action

Stop and restart the Device server and try the action again.

HWNPM2003E Initialization of the Device server event service failed. No performance threshold exception alerts will be generated.

Explanation

A failure occurred initializing the Device server event service. This means that no internal performance monitoring events can be externalized, including any threshold exception events that would normally cause the configured alerts to be generated. However, threshold exceptions are still logged and can be displayed using constraints

violation reports.

Action

Restart the Device Server service. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2004E Initialization of the product scheduler status service failed. The status of performance monitors will not be updated in the GUI.

Explanation

A failure occurred initializing the Device Server scheduler status service. This means that the Device server cannot update the status of any performance monitor jobs to be visible in the GUI. However, the monitors will continue to run normally, collecting performance data for their associated devices.

Action

Restart the Device Server service. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2005E Initialization of the product configuration data service failed. Performance monitors cannot be started without this service.

Explanation

A failure occurred while initializing the configuration data service. No new performance monitors can be started, and restarting prior monitors will fail.

Action

Restart the Device Server service. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2006E Initialization of the product configuration data service failed. Performance monitors cannot be started without this service.

Explanation

A failure occurred while initializing the performance data service. No new performance data collected from the devices can be saved in the database. Normally, performance monitors will terminate under these circumstances.

Action

Restart the Device server service. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2007E Initialization of the product counter data service failed for device type using agent type. Performance monitors will not be able to collect performance data from devices of this type.

Explanation

A failure occurred while initializing the counter data service. No new performance data can be collected from the devices of this type. This might be a transient failure due to the state of the product environment.

Action

Retry the operation. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2008E Initialization of the product metadata service failed. Performance monitors cannot be started without this service.

Explanation

A failure occurred while initializing the metadata service. No new performance monitors can be started, and restarting prior monitors will fail.

Action

Restart the Device server service. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2009E Unable to instantiate lower level service *service_class_name*.

Explanation

A failure occurred while trying to instantiate a particular lower level product service. This usually indicates an internal error.

Action

If you have made changes to the [TPC_installation_directory]\device\conf\pm.conf configuration file, reverse those changes. Restart the Device server service. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2010E Unable to instantiate the collection logic *implementation service class name*.

Explanation

A failure occurred trying to instantiate a particular implementation class in the Performance Manager. This usually indicates an internal error.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM2011E Unable to instantiate the performance statistics data class *class name*.

Explanation

A failure occurred trying to instantiate a particular data class in the Performance Manager. This usually indicates an internal error.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM2012I The product is using trace log directory log directory name .

Explanation

The specified log directory is in use.

Action

No action is required.

HWNPM2020W The performance monitor for device *device name* is not currently active, so a dynamic update of its monitor policy is not necessary.

Explanation

An attempt was made to dynamically update the monitor policy for a particular device, but the performance monitor for that device was not active. Policies have to be dynamically updated only for running monitors; for all other monitors updates are static, and take effect when the monitor for that device is started again.

Action

The policy updates are saved in the database, and will take effect when the performance monitor for the device is started. There is no effect on any other running monitors.

HWNPM2021W The performance monitor for device *device name* is not currently active, so a dynamic update of its threshold policy is not necessary.

Explanation

An attempt was made to dynamically update the threshold policy for a particular device, but the performance monitor for that device was not active. Policies have to be dynamically updated only for running monitors; for all other monitors updates are static, and take effect when the monitor for that device is started again.

Action

The policy updates are saved in the database, and will take effect when the performance monitor for the device is started. There is no effect on any other running monitors.

HWNPM2022E A performance monitor for device *device name* is already active. A new monitor for the same device cannot be started until the previous monitor completes or is cancelled.

Explanation

An attempt was made to start a new performance monitor for the specific device. There can be only a single performance monitor active within the IBM Spectrum Control environment for an individual device. Attempts to start a new monitor while a previous one is still running will fail.

Action

Cancel the previous monitor for the device before starting a new one, if so desired.

HWNPM2023W The performance monitor for device *device_name* is not currently active.

Explanation

An attempt was made to stop an existing performance monitor for the device. However, the performance monitor for that device was not active.

Action

No action is required.

HWNPM2024E Unable to find a monitor policy applicable to resource *resource name*.

Explanation

While starting a performance monitor for the specific resource, the monitor policy for the resource could not be determined. This error usually indicates a database access problem, because a policy is always in effect for every resource, even if it is only the default policy. However, a new performance monitor cannot be started for the resource until the problem is resolved.

Action

Try the following actions:

- Verify that the local area network is available.
- Verify that that the database repository is up and running
- Verify that the related database service is active.
- Start the performance monitor again.

HWNPM2025E Unable to find a threshold policy applicable to resource *resource_name*.

Explanation

While starting a performance monitor for the specific resource, the threshold policy for the resource could not be determined. This error usually indicates a database access problem, because a policy is always in effect for every resource, even if it is only the default policy. However, a new performance monitor cannot be started for the resource until the problem is resolved.

Action

Try the following actions:

- Verify that the local area network is available.
- Verify that that the database repository is up and running
- Verify that the related database service is active.
- Start the performance monitor again.

HWNPM2026I The performance monitor's primary process has failed unexpectedly. Attempting to recover from the failure.

Explanation

During the operation of the performance monitor a failure occurred which caused the monitor's primary process to be terminated. The Performance Manager will attempt to automatically recover from the failure by restarting the failed process. Performance data collection will be interrupted temporarily until the process has been successfully restarted.

Action

No action is required. To view information about the failure, view the trace logs. Additional messages will be issued to indicate the success or failure of the restart attempt.

HWNPM2027I The performance monitor threshold checker has failed unexpectedly. Attempting to recover from the failure.

Explanation

During the operation of the performance monitor a failure occurred which caused the monitor's threshold checker to be terminated. The Performance Manager engine will attempt to automatically recover from the failure by restarting the failed threshold checker. No thresholds will be generated for the performance data collected in the current interval.

Action

No action is required. For additional information, check the trace logs. Additional messages will be issued to indicate the success or failure of the attempt to restart the threshold checker.

HWNPM2028I The performance monitor purge process has failed unexpectedly. Attempting to recover from the failure.

Explanation

During the operation of the performance monitor a failure occurred which caused the monitor's purge processor to be terminated. The Performance Manager engine will attempt to automatically recover from the failure by restarting the failed purge processor. The purging of expired performance data will be interrupted temporarily until the process has been successfully restarted.

Action

No action required. For additional information, see the trace logs. Additional messages will be issued to indicate the success or failure of the recovery attempt.

HWNPM2029I Successfully recovered from the performance monitor failure.

Explanation

The automatic recovery attempt of the earlier failure was successful. The subject process has been restarted, and the performance monitor will continue to operate.

Action

None.

HWNPM2030E Unable to recover from the performance monitor failure. The performance monitor for the storage resource will be shut down.

Explanation

The attempt to restart the performance monitor failed.

Action

Restart the performance monitor. For additional information regarding the cause of the original failure, or the failure to restart the performance monitor, see the trace logs. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

HWNPM2031E The performance monitor failed due to an internal error.

Explanation

If the error is recoverable, the performance manager will attempt to restart. This failure might be temporary, and will result in a temporary disruption of performance data collection for the device.

Action

Restart the performance monitor if it does not restart automatically. For additional information regarding the cause of the original failure, or the failure to restart the performance monitor, see the trace logs. If the problem persists, contact IBM Support.

Related reference

- [Getting support](#)

HWNPM2032W The performance monitor for device *device name* is not currently using the default monitor policy, so a dynamic update of the policy is not necessary.

Explanation

An attempt was made to dynamically update the default monitor policy for a particular device, but the performance monitor for that device is using an already customized monitor policy rather than the default policy. The update is ignored for this device.

Action

The new default monitor policy may have an affect on the running performance monitors of other devices, but it will not affect the running performance monitor of this device. To modify the behavior of the performance monitor of this device, please update the custom monitor policy in effect for this device.

HWNPM2033W The performance monitor for device *device name* is not currently using the default threshold policy, so a dynamic update of the policy is not necessary.

Explanation

An attempt was made to dynamically update the default threshold policy for a particular device, but the performance monitor for that device is using an already customized threshold policy rather than the default policy. The update is ignored for this device.

Action

The new default threshold policy may have an affect on the running performance monitors of other devices, but it will not affect the running performance monitor of this device. To modify the behavior of the performance monitor of this device, please update the custom threshold policy in effect for this device.

HWNPM2040E The device key specified for the snapshot vote (*key*) was not found in the database. The device does not exist.

Explanation

During the internal operation of the product, an invalid device ID was used.

Action

This issue resolves without user action. If these error messages persist, contact your IBM support representative. If this error message persists, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2050E Failed to get the latest configuration data for device *device_name*.

Explanation

There was an error either in getting a reference for the Device server configuration data service, or in retrieving the latest configuration data for the device using the reference to the configuration data service.

Action

Try to run a probe for the device. If the problem continues, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2051E No performance data was collected from device *device_name* for the current collection interval (*time_stamp*) because the performance monitor was stopped.

Explanation

The performance monitor was stopped while it was collecting data from the device.

Action

No action is required.

HWNPM2052E No performance data was collected from device *device_name* for the current collection interval due to an error. Data was last collected at *time_stamp*.

Explanation

No performance data was received from the device for the current collection interval. This could be due to an internal error, or a problem communicating with the device or device agent.

Action

No action is required. The current performance monitor is not affected and will continue to try to collect performance data for the device.

HWNPM2053E The new performance data collected from device *device_name* could not be saved in the database. Increase the size of the transaction log.

Explanation

The transaction log size of the database is too low. This is preventing large transactions such as performance data collections from being saved to the database. The current set of performance data could not be saved.

Action

Increase the size of your database transaction log and try the operation again.

HWNPM2054E The new performance data collected from device *device_name* could not be saved in the database. Increase the size of the database lock list.

Explanation

The lock list size of the database is too low. This is preventing large transactions such as performance data collections from being saved to the database. The current set of performance data could not be saved.

Action

Increase the size of your database lock list and try the operation again.

HWNPM2055E The new performance data collected from resource *device name* could not be saved.

Explanation

An error occurred when the product attempted to save performance data about the resource.

Action

Try the following actions:

- Verify that the local area network is available.
- Check the status of the Device server on the Home > System Management page.
- Check for error messages in the log file of the Device server.
- Verify that the database repository is up and running.
- Verify that the related database service is active.
- Try the action again.

For more information about troubleshooting, go to IBM Documentation at <https://www.ibm.com/docs/en/spectrum-control>.

HWNPM2056E No performance data was collected from device *device name* for the current performance monitor job duration. The performance monitor job status is set to 'failed'.

Explanation

No performance data was received from the device for the current job duration. This could be due to the fact that there are no IOs on the device.

Action

None. The running performance monitor is not affected, and will continue to try to collect performance data for the device until the duration has elapsed.

HWNPM2057E No performance data was collected from device *device_name* for the current collection interval because the performance monitor was stopped.

Explanation

The performance monitor was stopped while it was collecting data from the device.

Action

No action is required.

HWNPM2058E No performance data was collected from device *device_name* for the current collection interval due to an error.

Explanation

No performance data was received from the device for the current collection interval. This could be due to an internal error, or a problem communicating with the device or device agent.

Action

No action is required. The current performance monitor is not affected and will continue to try to collect performance data for the device.

HWNPM2060W The device does not support performance management for segment pool *pool ID*. Only incomplete performance data can be collected for array *array ID*.

Explanation

The specified segment pool contains multiple ranks, which makes it impossible to accurately manage the performance for those ranks, the arrays associated with those ranks, and the device adapters associated with those arrays.

For DS6000 and DS8000 devices whenever a segment pool contains multiple ranks, any volumes allocated in that segment pool might be spread across those ranks in an unpredictable manner. This makes it impossible to determine the performance impact of the volumes on the individual ranks. To avoid presenting the user with inaccurate or misleading performance data, the Performance Manager does not attempt to compute the performance metrics for the affected arrays and device adapters.

Action

None.

HWNPM2061W The device does not support performance management for segment pool *pool ID*. Only incomplete performance data can be collected for device adapter *DA ID*.

Explanation

The specified segment pool contains multiple ranks, which makes it impossible to accurately manage the performance for those ranks, the arrays associated with those ranks, and the device adapters associated with those arrays.

For DS6000 and DS8000 devices whenever a segment pool contains multiple ranks, any volumes allocated in that segment pool might be spread across those ranks in an unpredictable manner. This makes it impossible to determine the performance impact of the volumes on the individual ranks. To avoid presenting the user with inaccurate or misleading performance data, the Performance Manager does not attempt to compute the performance metrics for the affected arrays and device adapters.

Action

None.

HWNPM2062W Invalid error message saved in database

Explanation

The program tried to save an error message in the database, but that message was invalid.

Action

None.

HWNPM2100E The performance monitor for resource *device name* cannot be started because configuration data for the resource is not available.

Explanation

When a performance monitor was started for the resource, configuration data for the resource could not be retrieved, or configuration data for historical purposes could not be copied.

Action

Try the following actions:

- Verify that the local area network is available.
- Check the status of the Device server on the Home > System Management page.
- Check for error messages in the log file of the Device server.
- Verify that that the database repository is up and running.
- Verify that the related database service is active.
- Try the action again.

For more information about troubleshooting, go to IBM Documentation at <https://www.ibm.com/docs/en/spectrum-control>.

HWNPM2101E All agents that can collect performance data for resource *device name* are currently non-operational.

Explanation

A performance monitor could not collect performance data with the agents that are configured for the resource. When this error occurs, the performance monitor will be stopped or it will retry the agents periodically until an operational agent is found. You can determine which action is taken by selecting an option for the Sampling Failed alert in the definition of the performance monitor.

Action

To resolve the problem, try the following actions:

- Verify that the local area network is available.
- Verify that that the agents for the resource are up and available.
- Verify that the product can communicate with at least one of the agents.
- Verify that the resource is up and available.
- Start the performance monitor again.

HWNPM2102E The performance monitor for resource *device name* cannot be started because the resource might not support the collection of performance data.

Explanation

When a performance monitor starts, the product verifies that the monitored resource supports the collection of performance data. If performance support cannot be verified, the following problems might exist: the resource does not support performance data collection, older microcode exists on the resource, or the associated SMI-S provider is an older version.

Action

If you believe that the resource does support performance data collection, ensure that the proper microcode levels are installed on that resource and the associated SMI-S provider is up to date. If the problem persists, contact the vendor of the resource.

HWNPM2103W Agent *agent name* is non-operational. Attempting to find an alternative agent.

Explanation

The agent that has been used to retrieve performance statistics from the resource is no longer returning statistics. Contact with a different agent will be attempted, if one is defined. There will be additional messages issued to notify you whether the attempted change in agents was successful.

Action

Ensure that the specified agent is up and running. Respond to any further messages.

HWNPM2104I The performance monitor policy was adjusted due to agent limitations. Current values in effect are: *interval-length=interval-length, frequency=frequency*.

Explanation

After switching to an alternative agent, it may be necessary to adjust the performance manager policy because different types of agents may have different capabilities. If the new agent does not support as low an interval-length or frequency as the previous agent, this message is issued, and the policy that is in effect is temporarily adjusted as indicated in the message.

Action

None. The policy in effect will return to the original user-specified values after the original agent is reestablished for performance data collection for the monitor.

HWNPM2105E The performance monitor for resource *resource name* failed because the resource for enabling performance data collection cannot be reached.

Explanation

When starting a performance monitor for the resource, a connection to the resource or associated agent could not be established. This problem might occur when there is a network problem, or there is an undiscovered change in the network attributes of the resource or associated agent, such as a change to an IP address or port number. It might also occur when the resource or associated agent is down or not available.

Action

If the networking attributes of the resource or associated agent have changed, run a probe. Ensure that the proper user name and password is being used, and (if applicable) certificate files were specified for communication with the resource or agent.

Ensure that a network path exists to resource or agent, including appropriate pass-through of any firewalls between the two network end-points. If you suspect network connectivity problems, contact your network administrator.

Ensure that the resource or agent is operational. Check the log files to verify that no errors exist that would prevent it from accepting connection requests.

HWNPM2106E The performance monitor for device *device name* failed because of errors trying to enable performance data collection on the device or device agent: *error description*

Explanation

While starting a performance monitor for the specific device, the Performance Manager was able to establish a connection to the device or device agent, but was unable to initiate performance data collection on the device and/or device agent.

Action

Ensure that the device and/or device agent are fully operational. Check for error condition in the logs of the device or device agent.

HWNPM2107E The performance monitor for device *device name* failed because of unrecognized errors trying to enable performance data collection on the device or device agent: *error description*

Explanation

While starting a performance monitor for the specific device, the Performance Manager was able to establish a connection to the device or device agent, but was unable to initiate performance data collection on the device and/or device agent. This might be due to an internal error in the IBM Spectrum Control Performance Manager.

Action

Ensure that the device and/or device agent are fully operational. Retry the operation. If the problem persists, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM2108E The performance monitor for resource *resource name* failed during shutdown because the resource cannot be reached for terminating data collection.

Explanation

When stopping a performance monitor for the resource, a connection to the resource or associated agent could not be established. This problem might occur when there is a network problem, or there is an undiscovered change in the network attributes of the resource or associated agent, such as a change to an IP address or port number. It might also occur when the resource or associated agent is down or not available.

Action

If the networking attributes of the resource or associated agent have changed, run a probe. Ensure that the proper user name and password is being used, and (if applicable) certificate files were specified for communication with the resource or agent.

Ensure that a network path exists to resource or agent, including appropriate pass-through of any firewalls between the two network end-points. If you suspect network connectivity problems, contact your network administrator.

Ensure that the resource or agent is operational. Check the log files to verify that no errors exist that would prevent it from accepting connection requests.

Because the performance monitor is shutting down anyway, there is no immediate affect on operations. However, some resources or agents might not allow a subsequent restart of performance data collection, unless the previous collection was terminated properly. For some resources, leaving performance data collection activated can cause additional burden when processing their related resources. In either case, you can also try to manually stop the performance data collection on the resource by using the interface for that resource.

HWNPM2109E The performance monitor for resource *resource name* failed during shutdown because of errors during termination of performance data collection: *error description*

Explanation

When stopping a performance monitor for the resource, a connection was established, but the process for collecting performance data could not be terminated. This problem might occur when there is a network problem, or there is an undiscovered change in the network attributes of the resource or associated agent, such as a change to an IP address or port number. It might also occur when the resource or associated agent is down or not available.

Action

Because the performance monitor is shutting down anyway, there is no immediate affect on operations. However, some resources or agents might not allow a subsequent restart of performance data collection, unless the previous collection was terminated properly. For some resources, leaving performance data collection activated can cause additional burden when processing their related resources. In either case, ensure that the resource or agent are fully operational. Check for error condition in the logs of the resource or agent. You can also try to manually stop the performance data collection on the resource using the interface for that resource.

HWNPM2110E The performance monitor for resource *resource name* failed during shutdown because of unrecognized errors during termination of performance data collection: *error description*

Explanation

When stopping a performance monitor for the resource, a connection was established, but the process for collecting performance data could not be terminated. This might be due to an internal error in the product.

Action

Because the performance monitor is shutting down anyway, there is no immediate affect on operations. However, some resources or agents might not allow a subsequent restart of performance data collection, unless the previous collection was terminated properly. For some resources, leaving performance data collection activated can cause additional burden when processing their related resources. In either case, try to manually stop the performance data collection on the resource by using the interface for that resource. If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2111E The performance monitor for resource *resource name* failed because of errors retrieving the most recent configuration data for the resource.

Explanation

When starting a performance monitor for the resource, its configuration data could not be retrieved and processed. This problem might occur when there is corrupt or invalid configuration data in the database repository.

Action

Try running a probe job for the resource and starting the performance monitor again. If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2112I Agent *agent name* was selected for performance data collection from resource *resource name*.

Explanation

This message shows how performance data was collected for the resource. Performance data can be collected either through an SMI-S provider or a direct connection. This message is useful in environments where there are multiple agents that manage the same resource.

Action

No further action is required.

HWNPM2113I The performance monitor for resource *resource name* is starting in an active state.

Explanation

Performance monitors start in an active state when initiated due to a user or a scheduled action, or due to a server restart when the monitor was previously active and has not reached its intended duration. The monitor will collect performance data from the resource when its initialization has completed.

Action

No further action is required.

HWNPM2114I The performance monitor for resource *resource name* is starting in a dormant state.

Explanation

Performance monitors start in a dormant state due to a server restart, if the monitor was previously dormant or if it was previously active and has reached its intended duration. A dormant monitor does not collect new performance data from its associated resource, but only performs basic maintenance functions such as summarization and purge processing.

Action

No further action is required. If you want to reactivate the performance monitor, schedule or run it immediately.

HWNPM2115I Monitor Policy: name=*policy name*, creator=*policy creator*, description=*policy description*

Explanation

This message provides information regarding a particular performance monitor policy that is in effect for a particular device.

Action

None.

HWNPM2116I Monitor Policy: retention period: sample data=*length in days* days, hourly data=*length in days* days, daily data=*length in days* days.

Explanation

This message provides information about a performance monitor policy that is in effect for a resource.

Action

No further action is required.

HWNPM2117I Monitor Policy: interval length=*length in seconds* secs, frequency=*length in seconds* secs, duration=*length in hours* hours.

Explanation

This message provides information about a performance monitor policy that is in effect for a resource.

Action

No further action is required.

HWNPM2118I Threshold Policy: name=*policy name*, creator=*policy creator*, description=*policy description*

Explanation

This message provides information about a performance monitor threshold policy that is in effect for a resource.

Action

No further action is required.

HWNPM2119I Threshold Policy: retention period: exception data=*length in days* days.

Explanation

This message provides information about a performance monitor threshold policy that is in effect for a resource.

Action

No further action is required.

HWNP2120I Threshold Policy: *threshold name=name, component=component type, enabled=Yes or No, boundaries=critical stress boundary,warning stress boundary,warning idle boundary,critical idle boundary units.*

Explanation

This message provides information about a performance monitor threshold policy that is in effect for a resource.

Action

No further action is required.

HWNP2121I Monitor Policy: *interval length=length in seconds secs, frequency=length in seconds secs, duration=continue indefinitely.*

Explanation

This message provides information about a performance monitor policy that is in effect for a resource.

Action

No further action is required.

HWNP2122W No valid performance data was provided by the monitored resource. No performance data records were inserted into the database.

Explanation

The performance monitor contacted the resource and tried to collect data. However, the resource did not provide valid performance counter information.

If message HWNP2124W is also displayed, the resource was able to provide performance data, but it was determined to be invalid and discarded.

Action

The operation of the performance monitor is not affected, and it will attempt to retrieve performance data again for the next sample interval. If the monitored resource continues to provide no performance data, ensure that it is fully operational. If appropriate, also ensure that performance functionality is enabled for the resource.

HWNP2123I Performance data for resource *timestamp date and time* was collected and processed successfully. *record count* performance data records were inserted into the database repository.

Explanation

This message provides information for an active performance monitor. The performance data was collected at the indicated resource time in the server timezone. The indicated number of performance statistics records were saved in the database repository. The saved information was either received from the resource, or was computed based on the information received from the resource.

Action

None.

HWNPM2124W Performance data continuity is broken. The device was possibly reset or rebooted. record count performance data records were discarded.

Explanation

The message indicates that invalid performance information was received from the device.

In general, performance information is represented as a set of ever-increasing counters, and actual statistics are computed by taking the difference between two consecutive sets of such counters.

However, if a counter appears to decrease rather than increase between consecutive sets, the information is unusable and is discarded.

Counters can be expected to decrease if they are reset to zeros, which might happen normally when a device is reset or rebooted. For example, when new firmware is loaded, or in some cases when a device agent (such as an SMI-S provider) is reset or rebooted. In those situations, this warning message can be safely ignored.

If this warning is displayed when no reset or restart happened, the device or device agent might be generating incorrect performance statistics. You might have to contact your device vendor for further instructions.

Action

Determine whether the device or device agent was reset or rebooted.

Those situations include loading of new firmware, or fail-over and fail-back scenarios for Enterprise Storage Server, DS6000, and DS8000 storage subsystems.

In any of these cases, the reset of performance counters is expected behavior, and this warning message can be safely ignored. If it is not one of these cases, the device might be generating incorrect performance data, which might or might not lead to inaccurate performance reports.

Contact your device vendor for further instructions in those cases.

HWNPM2125W Aggregated performance values have been computed from the remaining data records, but their accuracy cannot be guaranteed.

Explanation

This message only appears in combination with message HWNPM2124W. See the description of this message for background information on the condition that has occurred. Configuration parameter PM.LimitCheckLenient controls whether or not any data will be saved in the IBM Spectrum Control database when HWNPM2124W occurs. If set to true, then only those performance records where invalid counters were specifically detected are discarded (indicated with message HWNPM2124W), but all other data for that sample interval is saved (indicated with message HWNPM2123I). However because a number of records had to be discarded, any performance statistics derived from the remaining data (for example the aggregated statistics for the entire subsystem) will be inaccurate because they will not include the activity represented by that discarded data.

Action

Use any performance data saved for this sample interval cautiously, due to its inherent inaccuracy. To avoid such potentially inaccurate values from being saved in the IBM Spectrum Control database, you can set the PM.LimitCheckLenient configuration parameter back to false by running the following command:

```
setdscfg -property PM.LimitCheckLenient -context PerformanceManager false
```

However this will mean that no performance data will be saved in the database for any sample intervals where invalid counters are detected.

Related reference

- <https://www.ibm.com/docs/en/spectrum-control/latest?topic=provisioning-changing-default-host-definition-provisioned-storage>

HWNPM2126I The performance monitor for device *device name* is stopping because its intended duration has elapsed.

Explanation

Performance monitors with a preset duration end when that duration has elapsed. The monitor job might show normal completion or failure, depending on the success or failure of the data collection process during its life.

Action

None. If you want to reactivate the monitor, start another monitor job for the device.

HWNPM2127I The performance monitor for device *device name* is stopping due to a user request.

Explanation

The performance monitor is ending because a user requested the monitor to stop. A user request can be submitted either from the IBM Spectrum Control graphical user interface, or the IBM Spectrum Control command line interface.

Action

None. If you want to reactivate the monitor, start another monitor job for the device.

HWNPM2128E The performance monitor for device *device name* is stopping due to an unexpected failure.

Explanation

The performance monitor is ending due to an environmental condition or due to an internal error. Additional messages prior to this message indicate the cause of the failure.

Action

In case of an environmental condition, such as a network problem, correct the problem and restart the monitor. In case of an internal error, contact your service representative. For additional details regarding the cause of the failure, see the trace logs.

Related reference

- [Getting support](#)
- [Default locations of log files](#)

HWNPM2129I The performance monitor for device *device name* is stopping because of a shutdown request.

Explanation

The performance monitor is ending because the IBM Spectrum Control Performance Manager Service is shutting down, usually due to the IBM Spectrum Control device server being stopped.

Action

None. When the Performance Manager Service is restarted, the previously running monitors will be automatically restarted, as long as their originally intended duration has not been exceeded.

HWNPM2130W Failed to retrieve the latest configuration data for device *device name*.

Explanation

During the normal operation of a performance monitor, it will periodically attempt to retrieve the latest configuration data for its associated device. This attempt failed for the monitor of the specified device.

The immediate operation of the running performance monitor is unaffected. However if the failure persists for prolonged periods of time, it is possible that a discrepancy between the saved performance data collected from the device and the configuration data of that device could develop. In that case, the performance reports or any advanced analysis of the performance data might yield misleading results.

Action

Try running a probe job for the specific device, and retry starting the performance monitor. If the problem persists, contact your IBM support representative.

HWNPM2131W Performance data could not be collected for device *device name*, because the device or data source cannot be reached (reason reason code). The current samples are skipped.

Explanation

The current attempt to retrieve a set of performance data from the device failed. No performance data will be inserted for the device in this time period. The next performance data sample recorded into the database might represent an average over more than the configured interval length.

The immediate operation of the running performance monitor is unaffected. The reason code can be used to help identify the exact cause of the problem encountered:

1. Reason Code 0 indicates that the exact reason for the failure could not be determined. This should be a rare occurrence.
2. Reason Code 1 indicates a bad target (device or data source) address. This condition can occur when the user-specified host name or IP address, or the target port number are invalid such that they would cause the formation of an invalid URL or IP Address. This type of failure is rare and can usually also be identified via a `java.net.MalformedURLException` printed in the trace logs.
3. Reason Code 2 indicates the problem to be an unknown target address. This condition can occur when a hostname rather than an IP address was specified as target address, and:
 - either the network is down,
 - the specified hostname cannot be resolved (i.e. the nameserver cannot be contacted, or the nameserver is down, or the specified hostname is not known to the nameserver), or
 - the specified hostname can be resolved by the nameserver but no longer exists on the network.This type of failure can usually also be identified via a `java.net.UnknownHostException` printed in the trace logs.
4. Reason Code 3 indicates the problem to be an unreachable target address. This condition can occur when an IP address rather than a hostname was specified as target address, and either the network or a part of the network is down or is blocked by a firewall (the host cannot be contacted), or the specified IP address does not exist on the network. This type of failure can usually also be identified via a `java.net.NoRouteToHostException` in the trace logs.
5. Reason Code 4 indicates the problem to be an unresponsive target. This condition can occur when the target server is powered off, or when the server is not listening on the port which is the target of the communication (for example if the web server or SMI-S provider is not operational). This type of failure can usually also be identified via a `java.net.ConnectException` printed in the trace logs.
6. Reason Code 5 indicates a communication time-out for communication that uses UDP rather than TCP, for example when using SNMP data sources. This condition can occur when the target server cannot be reached, or when the SNMP data source is disabled on the target server, or when the SNMP port (161) is blocked by a firewall.

Action

If the problem persists for an hour or longer, ensure that your device and data source (if applicable) are operational. Also ensure that a network path exists between the IBM Spectrum Control server and the device or data source, including any appropriate firewall pass-throughs. If there is nothing wrong with the device or device agent, or with the network path, try cancelling and restarting the performance monitor job. If the problem is still not resolved, contact your IBM support representative.

HWNPM2132W Performance data could not be collected for device *device name*. The current samples are skipped. (error description)

Explanation

The current attempt to retrieve a set of performance data from the device failed. No performance data will be inserted for the device in this time period. The immediate operation of the running performance monitor is unaffected. However the next performance data sample recorded into the database might represent an average over more than the configured interval length.

Action

If the problem persists for an hour or longer, ensure that your device and device agent (if applicable) are operational, and that performance data collection is still enabled. Try cancelling and restarting the performance monitor job. If the problem is still not resolved, contact your IBM support representative.

For a NetApp device it is recommended to put all the NetApp volumes online before collecting performance data.

HWNPM2133W Performance data could not be collected for device *device name* due to an unknown error. The current samples are skipped.

Explanation

The current attempt to retrieve a set of performance data from the device failed. No performance data will be inserted for the device in this time period. The immediate operation of the running performance monitor is unaffected. However the next performance data sample recorded into the database might represent an average over more than the configured interval length.

Action

If the problem persists for an hour or longer, ensure that your device and device agent (if applicable) are operational, and that performance data collection is still enabled. Try cancelling and restarting the performance monitor job. If the problem is still not resolved, contact your IBM support representative.

HWNPM2134W The state of the performance monitor for resource *resource name* started, but the status of the performance monitor was not updated.

Explanation

The status of the performance monitor was not updated and won't be shown in the GUI.

Action

Ensure that the product database is up and running. If the problem persists, contact IBM Software Support.

HWNPM2135W The state of the performance monitor for device *device name* has changed to 'active', but could not be recorded appropriately.

Explanation

The state of the specified monitor has changed to 'active'. However the IBM Spectrum Control user interfaces might not be able to display this fact due to a failure in updating the monitor's state record in the database.

Action

This error is usually indicative of a database problem. Ensure that the IBM Spectrum Control database is operational. If the problem persists, contact your IBM support representative.

HWNPM2136W The performance monitor for the resource *resource name* generated a warning, but the status of the performance monitor was not updated.

Explanation

The status of the performance monitor was not updated and won't be shown in the GUI.

Action

Ensure that the product database is up and running. If the problem persists, contact IBM Software Support.

HWNPM2137W The performance monitor for the resource *resource name* stopped, but the status of the performance monitor was not updated.

Explanation

The status of the performance monitor was not updated and won't be shown in the GUI.

Action

Ensure that the product database is up and running. If the problem persists, contact IBM Software Support.

HWNPM2138W The performance monitor for the resource *resource name* completed the collection of data, but the status of the

performance monitor was not updated.

Explanation

The status of the performance monitor was not updated and won't be shown in the GUI.

Action

Ensure that the product database is up and running. If the problem persists, contact IBM Software Support.

HWNPM2139W The performance monitor for the resource *resource name* failed, but the status of the performance monitor was not updated.

Explanation

The status of the performance monitor was not updated and won't be shown in the GUI.

Action

Ensure that the product database is up and running. If the problem persists, contact IBM Software Support.

HWNPM2140W The status of the performance monitor for the resource *resource name* was not updated.

Explanation

The status of the performance monitor was not updated and won't be shown in the GUI.

Action

Ensure that the product database is up and running. If the problem persists, contact IBM Software Support.

HWNPM2141E The service is unavailable because an unexpected error occurred.

Explanation

The request for the performance service wasn't accepted.

Action

Contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2142E Performance data can't be collected for the resource *resource name* because the performance monitor was disabled.

Explanation

The performance monitor must remain enabled until the scheduled collection of the data can be completed. If the performance monitor is disabled by a user, application, or external entity while the performance monitor is running, the collection of performance data can't be completed.

Action

Determine whether another application or user has disabled the collection of performance data for the storage resource and restart the performance monitor.

If the resource or the agent for the resource, such as the SMI-S provider, is shut down or restarted, the performance monitor can be disabled. In such cases, the performance monitor can be restarted after the resource or its agent is restarted.

HWNPM2143E The performance monitor for the resource *resource name* was started, but the status of the performance monitor was not updated and might not be shown in the GUI.

Explanation

The current status of the performance monitor might not be shown in the GUI.

Action

Ensure that the product database is up and running. If the problem persists, contact IBM Software Support.

Related reference

- [Getting support](#)

HWNPM2144W The performance data cannot be checked against the alert conditions, so no alerts can be generated.

Explanation

After every new set of performance data is collected by the performance monitor, the data is checked to determine whether it violates any alert conditions. When violations are detected, the appropriate alerts are generated and any necessary user notifications are sent.

This message indicates that the normal checking against the new data cannot be done. Because the checking cannot be done, violations are not detected, alerts are not generated, and no notifications are sent.

Action

Go to Home > System Management to verify that the Alert Server component and the Database component are up and running. If these components are not operational, restart them. If the problem persists, contact IBM Software Support.

HWNPM2145I The data is being collected by the data collector: *data collector host*.

Explanation

This message is for informational purposes only.

Action

No action is required.

HWNPM2146W Performance data could not be collected for device *device name*, the exact reason for the failure could not be determined. The current samples are skipped.

Explanation

The current attempt to retrieve a set of performance data from the device failed. No performance data will be inserted for the device in this time period. The next performance data sample recorded into the database might represent an average over more than the configured interval length.

The immediate operation of the running performance monitor is unaffected.

Action

If the problem persists for an hour or longer, ensure that your device and device agent (if applicable) are operational. Also ensure that a network path exists between the TPC server and the device or device agent, including any appropriate firewall pass-throughs. If there is nothing wrong with the device or device agent, or with the network path, try cancelling and restarting the performance monitor job. If the problem is still not resolved, contact your IBM support representative.

HWNPM2147W Performance data could not be collected for device *device name*, because of a bad target (device or agent) address. The current samples are skipped.

Explanation

The current attempt to retrieve a set of performance data from the device failed. No performance data will be inserted for the device in this time period. The next performance data sample recorded into the database might represent an average over more than the configured interval length. This condition can occur when the user-specified host name or IP address, or the target port number are invalid such that they would cause the formation of an invalid URL or IP Address. This type of failure is rare and can usually also be identified via a `java.net.MalformedURLException` printed in the trace logs.

The immediate operation of the running performance monitor is unaffected.

Action

If the problem persists for an hour or longer, ensure that your device and device agent (if applicable) are operational. Also ensure that a network path exists between the TPC server and the device or device agent, including any appropriate firewall pass-throughs. If there is nothing wrong with the device or device agent, or with the network path, try cancelling and restarting the performance monitor job. If the problem is still not resolved, contact your IBM support representative.

HWNPM2148W Performance data could not be collected for device *device name*, because of an unknown target address. The current samples are skipped.

Explanation

The current attempt to retrieve a set of performance data from the device failed. No performance data will be inserted for the device in this time period. The next performance data sample recorded into the database might represent an average over more than the configured interval length. This condition can occur when a hostname rather than an IP address was specified as target address, and either the network is down, the specified hostname cannot be resolved (i.e. the nameserver cannot be contacted, or the nameserver is down, or the specified hostname is not known to the nameserver), or the specified hostname can be resolved by the nameserver but no longer exists on the network. This type of failure can usually also be identified via a `java.net.UnknownHostException` printed in the trace logs.

The immediate operation of the running performance monitor is unaffected.

Action

If the problem persists for an hour or longer, ensure that your device and device agent (if applicable) are operational. Also ensure that a network path exists between the TPC server and the device or device agent, including any appropriate firewall pass-throughs. If there is nothing wrong with the device or device agent, or with the network path, try cancelling and restarting the performance monitor job. If the problem is still not resolved, contact your IBM support representative.

HWNPM2149W Performance data could not be collected for device *device name*, because of an unreachable target address. The current samples are skipped.

Explanation

The current attempt to retrieve a set of performance data from the device failed. No performance data will be inserted for the device in this time period. The next performance data sample recorded into the database might represent an average over more than the configured interval length. This condition can occur when an IP address rather than a hostname was specified as target address, and either the network or a part of the network is down or is blocked by a firewall (the host cannot be contacted), or the specified IP address does not exist on the network. This type of failure can usually also be identified via a `java.net.NoRouteToHostException` in the trace logs.

The immediate operation of the running performance monitor is unaffected.

Action

If the problem persists for an hour or longer, ensure that your device and device agent (if applicable) are operational. Also ensure that a network path exists between the TPC server and the device or device agent, including any appropriate firewall pass-throughs. If there is nothing wrong with the device or device agent, or with the network path, try cancelling and restarting the performance monitor job. If the problem is still not resolved, contact your IBM support representative.

HWNPM2150W Performance data could not be collected for device *device name*, because of an unresponsive target. The current samples are skipped.

Explanation

The current attempt to retrieve a set of performance data from the device failed. No performance data will be inserted for the device in this time period. The next performance data sample recorded into the database might represent an average over more than the configured interval length. The problem is an unresponsive target. This condition can occur when the target server is powered off, or when the server is not listening on the port which is the target of the communication (for example if the web server or CIMOM is not operational). This type of failure can usually also be identified via a `java.net.ConnectException` printed in the trace logs.

The immediate operation of the running performance monitor is unaffected.

Action

If the problem persists for an hour or longer, ensure that your device and device agent (if applicable) are operational. Also ensure that a network path exists between the TPC server and the device or device agent, including any appropriate firewall pass-throughs. If there is nothing wrong with the device or device agent, or with the network path, try cancelling and restarting the performance monitor job. If the problem is still not resolved, contact your IBM support representative.

HWNPM2151W Performance data could not be collected for device *device name*, because a communication time-out for communication that uses UDP rather than TCP. The current samples are skipped.

Explanation

The current attempt to retrieve a set of performance data from the device failed. No performance data will be inserted for the device in this time period. The next performance data sample recorded into the database might represent an average over more than the configured interval length. This condition can occur when the target server cannot be reached, or when the SNMP data source is disabled on the target server, or when the SNMP port (161) is blocked by a firewall.

The immediate operation of the running performance monitor is unaffected.

Action

If the problem persists for an hour or longer, ensure that your device and device agent (if applicable) are operational. Also ensure that a network path exists between the TPC server and the device or device agent, including any appropriate firewall pass-throughs. If there is nothing wrong with the device or device agent, or with the network path, try cancelling and restarting the performance monitor job. If the problem is still not resolved, contact your IBM support representative.

HWNPM2200I The performance monitor successfully collected the configuration data for the storage system with the following internal resources: *number_of_pools* pools, *number_of_controllers* controllers, *number_of_device_adapters* device adapters, *number_of_ports* ports, *number_of_host_connections* host connections, *number_of_ranks* ranks, *number_of_arrays* arrays, and *number_of_volumes* volumes.

Explanation

The configuration data for the monitored storage system and its internal resources were collected.

The number of host connections that is shown here and in the log file might vary because only host connections with one or more assigned volumes are counted. To see the total number of host connections, go to the Overview page for the storage system.

Host connections with the same name and the same volume group are considered identical and are not double counted.

Action

No further action is required.

HWNPM2201I The performance monitor successfully collected the configuration data for the storage system with the following

internal resources: *number_of_io_groups* I/O Groups, *number_of_nodes* nodes, *number_of_ports* ports, *number_of_host_connections* host connections, *number_of_pools* pools, *number_of_managed_disks* managed disks, *number_of_local_disks* local disks, *number_of_volumes* volumes, and *number_of_volume_copies* volume copies.

Explanation

The configuration data for the monitored storage system and its internal resources were collected.

The number of host connections that is shown here and in the log file might vary because only host connections with one or more assigned volumes are counted. To see the total number of host connections, go to the Overview page for the storage system.

Host connections with the same name and the same volume group are considered identical and are not double counted.

The number of volume copies includes both the primary and secondary volume copies for all of the volumes.

Action

If the number of internal resources shown in the message doesn't match the number of actual internal resources for the storage system, run a probe on the storage system.

HWNPM2202I The performance monitor successfully retrieved the configuration data for the switch. The following internal resources were found: *number_of_trunks* trunks, and *number_of_ports* ports.

Explanation

The running performance monitor has updated its internal state with the latest configuration data from the switch. The specified internal resources were found.

Trunks are ISL Trunks, ICL Trunks, or port channels configured for the switch.

Action

None.

HWNPM2203I The performance monitor successfully retrieved the configuration data for the storage system. The following internal resources were found: *number_of_host_connections* host connections, *number_of_modules* modules, *number_of_ports* ports, *number_of_pools* pools, and *number_of_volumes* volumes.

Explanation

The running performance monitor has updated its internal state with the latest configuration data from the storage system. The specified internal resources were found.

Note: The number of host connections reflects only the host connections that have at least one volume assigned. Host connections with 0 volumes have no associated performance data.

In some cases, the number of host connections that is shown in a log file is less than the number of host connections that is shown on the Storage System details page in the web-based GUI. This difference might occur because log files do not count host connections that have no volumes assigned. However, the Storage System details page does count host connections with 0 volumes.

Action

None.

HWNP2204I The performance monitor successfully retrieved the configuration data for the storage system. The following internal resources were found: *number_of_nodes* nodes, *number_of_ports* ports, and *number_of_modules* flash modules.

Explanation

The running performance monitor has updated its internal state with the latest configuration data from the storage system. The specified internal resources were found.

Action

No action is required.

HWNP2205I The performance monitor successfully retrieved the configuration data for the storage system. The following internal resources were found: *number_of_ports* ports, *number_of_controllers* controllers, *number_of_volumes* volumes, and *number_of_disks* disks.

Explanation

The running performance monitor has updated its internal state with the latest configuration data from the storage system. The specified internal resources were found.

Action

No action is required.

HWNP3000E There was a problem establishing the database connection.

Explanation

An exception occurred when trying to retrieve the database connection from the connection pool.

Action

Make sure your database is working correctly. If you cannot find anything wrong with the database, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNP3001E An unexpected null row was returned from a database cursor.

Explanation

An unexpected null row was returned when trying to retrieve a row from a database cursor.

Action

Make sure your database is working correctly. If you cannot find anything wrong with the database, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM3002E An unexpected database exception occurred.

Explanation

An unexpected database exception occurred when trying to access the database.

Action

Make sure your database is working correctly. If you cannot find anything wrong with the database, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM3003E An unexpected database exception occurred on the snapshot database tables.

Explanation

An unexpected database exception occurred when trying to access the snapshot database tables.

Action

Make sure your database is working correctly. If you cannot find anything wrong with the database, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM3004E The snapshot ID could not be found.

Explanation

The snapshot ID could not be found in the snapshot database tables.

Action

Make sure your database is working correctly. If you cannot find anything wrong with the database, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM3500E The current transaction has been rolled back because of a deadlock.

Explanation

The application is rolled back to the previous COMMIT because of a deadlock.

Action

Contact your IBM Support Representative.

Related reference

- [Getting support](#)

HWNPM3501E The current transaction has been rolled back because of a timeout.

Explanation

The application is rolled back to the previous COMMIT because of a timeout.

Action

Check if there is an escalation problem with the database. If so, enlarge the database lock list.

HWNPM3502E The current transaction has been rolled back because the database transaction log has been exhausted.

Explanation

The application is rolled back to the previous COMMIT because the transaction log is full.

Action

Ensure there is sufficient disk space for the database transaction logs. Retry the operation and if the error persists, increase the DB transaction log size.

HWNPM3503E The current transaction has been rolled back because the database disk space has been exhausted.

Explanation

The application is rolled back to the previous COMMIT because there is insufficient disk space to write more data into the database.

Action

Increase the amount of disk space available for the database. If this is impossible, reduce your history retention settings to decrease the size of the database. If deleting data fails due to transaction logs also being exhausted, some tables or indices may have to be temporarily dropped. Contact your IBM Support Representative for assistance.

HWNPM3600E The threshold identifier parameter value : *threshold ID* is not valid.

Explanation

The Affected volumes and hosts Reporting feature is not supported for the threshold identifier parameter passed to the function.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM3601E The target component type parameter value : *component type* is not valid for the threshold identifier : *threshold ID* passed to the affected volumes and hosts reporting function.

Explanation

Each threshold identifier is associated with a target component type. The target component type parameter value passed in is not valid.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM3602E There was a problem retrieving the performance data needed to generate the affected volumes and hosts report for the device *device name*.

Explanation

An exception occurred when trying to retrieve the performance statistics data from the database that was needed to generate the affected volumes and hosts report.

Action

Ensure the IBM Spectrum Control database is operational and retry the operation. Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM3603E The sample volume performance data needed to generate the affected volumes and hosts report for the device *device name* was not found in the IBM Spectrum Control database.

Explanation

The sample volume performance statistics data collected during the time the threshold or constraint was violated is no longer present in the IBM Spectrum Control database. The affected volumes and hosts report cannot be generated in absence of this data.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM3604E There are no volumes associated with the specified target component, *component name*, in the IBM Spectrum Control database. Therefore, the resulting Affected Volumes and Hosts report will be empty.

Explanation

A constraint violation is always associated with a particular component (array, controller, I/O group, MDisk, etc.), which is the component that actually violated the constraint. In this case, the component associated with the constraint violation that was selected as the target for the Affected Volumes and Hosts report, currently has no volumes configured or assigned to it. This means that there are neither volumes nor hosts affected by this particular constraint violation, and the resulting report will be empty.

Action

It is possible that the IBM Spectrum Control database has outdated information for the corresponding device. If you believe that there are one or more volumes configured or assigned to the constraint violation's component, run a new probe of the device, to ensure that the latest configuration information is present in the IBM Spectrum Control database. Future constraint violations for this component (and for all other components of the device) will use this updated information, and should generate an accurate Affected Volumes and Hosts report.

HWNPM4000E Unable to retrieve the device agent that managed this device: *device identifier*.

Explanation

Unable to retrieve the device agent, such as an SMI-S provider, that controls this device. An invalid user ID, password, or namespace might be configured in IBM Spectrum Control for this device.

Action

Confirm that this device is registered with a device agent, and that the device agent is known to IBM Spectrum Control with correct user ID, password and namespace. If the device agent is correct, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4001E Timeout while starting performance data collection for this device: *device identifier*.

Explanation

An SMI-S provider communication timeout occurred while starting performance data collection.

Action

Increase the SMI-S provider communication timeout. If the problem persists, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4002E Unable to start performance data collection for this device: *device identifier*.

Explanation

No response was returned while communicating with this device or its device agent.

Action

Check that the device agent, or SMI-S provider, is functional and that the IBM Spectrum Control user ID, password, and namespace are correct for the device agent. If they are, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4003E Performance data collection has already been enabled for this device: *device identifier*.

Explanation

Another user has already started performance data collection, and this device or its device agent only allows one collection to be performed at a time.

Action

Wait for the previously enabled collection to complete, then try again.

Related reference

- [Getting support](#)

HWNPM4004E Failed to enable performance data collection for this device: *device identifier*.

Explanation

The device or device agent returned an error code when attempting to enable performance data collection.

Action

Check the frequency and duration of the performance data collection. If correct, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4005I Successfully enabled performance data collection on the storage subsystem, using device access point *SMI-S provider address*.

Explanation

Performance data collection was enabled successfully.

Action

None. While performance data collection is enabled, the storage subsystem will internally be generating statistics on the performance of various internal components, such as volumes, arrays, ports, and so forth. Performance data collection will be disabled again, when the user-specified duration of the performance monitor has elapsed.

HWNPM4006E An exception occurred while starting performance data collection for this device: *device identifier*.

Explanation

An exception occurred while attempting to start the performance data collection.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4007E A timeout occurred while stopping performance data collection for this device: *device identifier*.

Explanation

An SMI-S provider communication timeout occurred while stopping performance data collection.

Action

Increase the SMI-S provider communication timeout. If the problem persists, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4008E Unable to stop performance data collection for this device: *device identifier*.

Explanation

No response was returned while communicating with this device or its device agent.

Action

Check that the device agent, or SMI-S provider, is functional and that the IBM Spectrum Control user ID, password, and namespace are correct for the device agent. If they are, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4009E Performance data collection is not enabled for this device: *device identifier*.

Explanation

Performance data collection is not currently running for this device. It must be enabled before stopping it.

Action

None. Performance data collection is already stopped.

HWNPM4010E Failed to disable performance data collection for this device: *device identifier*.

Explanation

The device or device agent returned an error code when attempting to disable performance data collection.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4011I Successfully disabled performance data collection on the storage subsystem, using device access point *SMI-S provider address*.

Explanation

The performance data collection was disabled successfully.

Action

None. While performance data collection is disabled, the storage subsystem will not be generating statistics on the performance of its internal components. Performance data collection will be enabled again, when the next performance monitor is started for this device.

HWNPM4012E An exception occurred while stopping performance data collection for this device: *device identifier*.

Explanation

An exception occurred while attempting to stop the performance data collection.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4013E A timeout occurred while retrieving the status of the performance data collection for this device: *device identifier*.

Explanation

An SMI-S provider communication timeout occurred while retrieving the performance data collection's status.

Action

Increase the SMI-S provider communication timeout. If the problem persists, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4014E Unable to retrieve the status of the performance data collection for this device: *device identifier*.

Explanation

No response was returned while communicating with this device or its device agent.

Action

Check that the device agent, or SMI-S provider, is functional and that the IBM Spectrum Control user ID, password, and namespace are correct for the device agent. If they are, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4015I Performance data collection is not enabled for this device: *device identifier*.

Explanation

Performance data collection is not currently running for this device.

Action

None.

HWNPM4016I Performance data collection is enabled for this device: *device identifier*.

Explanation

Performance data collection is currently running for this device.

Action

None.

HWNPM4017E Unable to determine the status of the performance data collection for this device: *device identifier*.

Explanation

The device or device agent returned an error code when retrieving the status.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4018E Failed to retrieve the status of the performance data collection for this device: *device identifier*.

Explanation

An exception occurred while attempting to retrieve the performance data collection's status.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4019E A timeout occurred while polling the performance statistics for this device: *device identifier*.

Explanation

An SMI-S provider communication timeout occurred while polling the performance statistics.

Action

Increase the SMI-S provider communication timeout. If the problem persists, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4020E Unable to retrieve the performance statistics for this device: *device identifier*.

Explanation

No response was returned while communicating with this device or its device agent.

Action

Check that the device agent, or SMI-S provider, is functional and that the IBM Spectrum Control user ID, password, and namespace are correct for the device agent. If they are, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4021E No performance statistics available at the current time for this device: *device identifier*.

Explanation

No performance statistics available at the current time.

Action

The device agent might not be responding. Retry the performance data collection. If the problem persists, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4022E Failed to disable performance data collection for this device: *device identifier*.

Explanation

The device or device agent returned an error code when attempting to disable performance data collection.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4023W A set of performance statistics data was empty for this device: *device identifier*.

Explanation

At least one cluster from the specified device was down.

Action

None.

HWNPM4024E An exception occurred while stopping performance data collection for this device: *device identifier*.

Explanation

An exception occurred while attempting to stop the performance data collection.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4025E Unable to retrieve storage subsystem for this device: *device identifier*.

Explanation

The storage subsystem was unable to be retrieved from the database.

Action

Re-probe the storage subsystem.

HWNPM4026E Failed to retrieve storage subsystem for this device: *device identifier*.

Explanation

An Exception occurred when retrieving the subsystem information from the database.

Action

Re-probe the storage subsystem.

HWNPM4027E Failed to properly initialize counter data service for this device: *device identifier*.

Explanation

An Exception occurred when retrieving the subsystem information from the database.

Action

Check if the database contains the storage subsystem for this device.

HWNPM4028W Performance data cannot be collected because the *security role* authority of the user account *user name* for accessing *device identifier* is not sufficient.

Explanation

The user account that is used to log on to the storage system does not have the required authority for collecting performance data. For versions of IBM Spectrum Virtualize earlier than 8.3.1.2, the user must have the Administrator or SecurityAdmin role.

Action

If the resource is a Spectrum Virtualize storage system, complete one of the following tasks:

- Open the management GUI for the storage system and change the role of the user account to either 'Administrator' or 'SecurityAdmin'.
- In the IBM Spectrum Control GUI, go to the Block Storage Systems page, right-click the resource, and select Connections -> Modify Connection. Enter a different user name that has sufficient authority for collecting performance data. You might need to stop and restart the performance monitor to have these changes take effect. Right-click the resource again and select the appropriate Data Collection options.

HWNPM4029W Performance data cannot be collected because the collection of performance statistics is stopped on *device identifier*. The *security role* authority of the user account *user name* for accessing the storage system is not sufficient to start the collection of performance statistics.

Explanation

The user does not have sufficient authority to run the following command to start the collection:

```
svctask startstats -interval "time in minutes"
```

You might see this message if a user with a role other than Administrator, or SecurityAdmin is used to access the storage system.

Action

Complete one of the following tasks:

- Run the following command on the storage system to start the collection of performance statistics. You must be logged on as a user with the Administrator, or SecurityAdmin role.

```
svctask startstats -interval "time in minutes"
```

For example, `svctask startstats -interval 5`

The specified interval must be less than or equal to the performance monitor interval for the storage system in IBM Spectrum Control. To check the value, in the IBM Spectrum Control GUI, go to the Block Storage Systems page and view the value for Performance Monitor Interval for the storage system.

- Change the authority of the user that is used to log on to the storage system. Open the management GUI for the storage system and change the role of the user account to the Administrator, or SecurityAdmin role.

HWNPM4030W Performance data cannot be collected. The performance interval *device interval* on *device identifier* is greater than the sample interval and the set *security role* authority of the user account *user name* is not sufficient to update the interval value on *device identifier* .

Explanation

The interval that is specified on the storage system for the collection of performance statistics must be less than or equal to the performance monitor interval for the storage system in IBM Spectrum Control. The user that is used to access the storage system does not have sufficient authority to run the following command to change the interval:

```
svctask startstats -interval "time in minutes"
```

You might see this message if a user with a role other than Administrator, or SecurityAdmin is used to access the storage system.

Action

Complete one of the following tasks:

- Run the following command on the storage system to start the collection of performance statistics. You must be logged on as a user with the Administrator, or SecurityAdmin role.

```
svctask startstats -interval "time in minutes"
```

For example, `svctask startstats -interval 5`

The specified interval must be less than or equal to the performance monitor interval for the storage system in IBM Spectrum Control. To check the value, in the IBM Spectrum Control GUI, go to the Block Storage Systems page and view the value for Performance Monitor Interval for the storage system.

- Change the authority of the user that is used to log on to the storage system. Open the management GUI for the storage system and change the role of the user account to the Administrator, or SecurityAdmin role.

HWNPM4051E Failed to obtain a reference to the Performance Manager Configuration Data Service for this device: *device name*.

Explanation

The Switch counter data service was unable to access the configuration data service. Without the configuration data service, the Switch counter data service cannot run.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4052E Error occurred in trying to retrieve a device agent for this device: *device name*.

Explanation

An error occurred while attempting to obtain a device agent, such as an SMI-S provider, that controls this device.

Action

Confirm that this device is registered with a device agent, and that the device agent is known to IBM Spectrum Control. If there is a device agent, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4053E Unable to locate or retrieve the device agent that manages this device: *device name*.

Explanation

Unable to locate the device agent, such as an SMI-S provider, that controls this device, or an internal error occurred while attempting to retrieve the agent.

Action

Confirm that this device is registered with a device agent, and that the device agent is known to IBM Spectrum Control. If there is a device agent, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4054E Error occurred in trying to construct the poll state information for this device: *device name*.

Explanation

An error occurred while attempting to construct the poll state information, a prerequisite for collecting performance statistics.

Action

Re-probe the storage subsystem.

HWNPM4055E Unable to construct the poll state information for this device: *device name*.

Explanation

Unable to construct the poll state information for this device, or an internal error occurred while attempting to construct this information. This information is a prerequisite for collecting performance statistics.

Action

Re-probe the storage subsystem.

HWNPM4056E SMI-S provider operation triggered a timeout (step timeout= *step timeout value* seconds, operation timeout= *total timeout value* seconds,).

Explanation

The SMI-S provider communication timeout expired while executing an SMI-S provider operation for the switch device.

Action

Increase the SMI-S provider connection (individual step) timeout and communication (total operation) timeout for the device. If the problem persists, contact your IBM support representative.

Related reference

-  [Getting support](#)

HWNPM4057E Mismatch in device identifier for this device: *device name*.

Explanation

The Switch counter data service passed a different value of the unique device identifier for this operation from the value used earlier. The same value of the unique device identifier must be used for all invocations of Counter Data Service for Switch functionality. This is an error by the invoking code.

Action

Contact your IBM support representative.

Related reference

-  [Getting support](#)

HWNPM4058E Failed to build the parameter Map for this device: *device name*.

Explanation

Failed to build the parameter Map, which needs to be passed to the Discovery Service for collecting performance statistics. The Counter Data Service for the device cannot successfully collect performance statistics for this device without this initialization. This is an internal error.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4059I Performance data collection has already been enabled for this device: *device name*.

Explanation

Performance data collection has already been started for this device. It is good practice to stop the previously started collection before starting a fresh collection.

Action

None.

HWNPM4060I Performance data collection was successfully started for this device: *device name*.

Explanation

Performance data collection has been started for this device. It will continue until configured or stopped.

Action

None.

HWNPM4061E Performance data collection could not be started for this device: *device name*.

Explanation

Performance data collection could not be started for this device.

Action

Verify that the device is functioning and known to IBM Spectrum Control, supports the SMI-S Switch profile, and is reachable via a functioning access point that has been registered with IBM Spectrum Control. Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4062I Performance data collection successfully stopped for this device: *device name*.

Explanation

Performance data collection successfully stopped for this device.

Action

None.

HWNPM4063W Parse exception in performance data collected this device: *device name*.

Explanation

A parse error occurred while processing performance data collected for this device. An exception was caught while trying to extract an SMI-S provider Property from the data returned by the collector. The data is ignored.

Action

Contact your IBM support representative.

HWNPM4064E *wrong format in performance data collected for this device: device name.*

Explanation

The performance data object representing the performance statistics for a port of this device does not have the expected format. This is an internal error.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4065W *number of null time stamps null time stamp(s) for performance data collected from the device were substituted by server time stamp(s).*

Explanation

The switch device did not report the required SMI-S provider StatisticTime time stamp property for the reported number of statistics instances reported. These were substituted by IBM Spectrum Control server time stamps. If this is an intermittent problem, it might cause inconsistencies in performance reports. This is a device provider error.

Action

Contact your IBM support representative, and switch vendor.

HWNPM4066W *count of null operational status null Port Operational Status value(s) for performance data collected from the device was/were substituted by default value(s).*

Explanation

The switch device did not report the required SMI-S provider OperationalStatus property for the reported number of statistics instances reported. These were substituted by a default value of port operational status ('unknown'). This can be caused either because an assembled FCPort CIMInstance was used by Fabric Data collection code, or by an error in the provider for the device.

Action

If this warning occurred when collecting performance statistics for all the ports of a switch, contact your IBM support representative, and switch vendor.

HWNPM4081E *A database cursor operation failed.*

Explanation

A failure occurred when IBM Spectrum Control attempted to obtain or use a database cursor.

Action

Verify that the database is operational and online.

HWNPM4082E A database connect operation failed.

Explanation

A failure occurred when IBM Spectrum Control attempted to connect to the database.

Action

Verify that the database is operational and online.

HWNPM4083E A database retrieve operation failed.

Explanation

A failure occurred when IBM Spectrum Control attempted to retrieve a row from a database table. This is an internal error, that can occur either because of a problem with the database, or because an attempt was made to retrieve a non-existent row from the database.

Action

Verify that the database is operational and online. If it is verified to be so, contact your IBM support representative.

Related reference

-  [Getting support](#)

HWNPM4084E A database operation failed.

Explanation

A generalized database failure occurred.

Action

Verify that database is operational and online. If it is verified to be so, contact your IBM support representative.

Related reference

-  [Getting support](#)

HWNPM4085E A database query operation failed.

Explanation

A failure occurred when IBM Spectrum Control attempted a database query operation. This is an internal error, that can occur either because of a problem with the database, or because an attempt was made to retrieve a non-existent row from the database.

Action

Verify that database is operational and online. If it is verified to be so, contact your IBM support representative.

Related reference

-  [Getting support](#)

HWNPM4086W A database query gave no result rows.

Explanation

IBM Spectrum Control executed a query operation that generated no result rows. This might be because the parameters to the query were valid, but there is no underlying data, or the parameters to the query were incorrect. This might result in the failure of a higher level operation.

Action

Check the trace log for more information.

HWNPM4087W Missing or invalid association between SMI-S provider *SMI-S provider URL* and device *device name*. The configured SMI-S provider is inoperative, or may no longer be managing the specified device.

Explanation

The IBM Spectrum Control database no longer carries the association between the indicated SMI-S provider and device. This could be because a previous IBM Spectrum Control discovery job detected that the device was removed from the SMI-S provider configuration, or it could be because the SMI-S provider was not operational at the time of the last IBM Spectrum Control discovery job, causing the SMI-S provider-to-device association to be deleted from the IBM Spectrum Control database.

Action

There are several common causes for this problem:

1. If the device was moved to a different SMI-S provider, ensure that the most recent IBM Spectrum Control SMI-S provider discovery job detected the new SMI-S provider and its association to the device. If necessary, run a new SMI-S provider discovery job to detect this association. Then manually stop and restart the IBM Spectrum Control performance monitor for the device, to force the use of the new SMI-S provider by the monitor.
2. If the device is intended to still be managed by the indicated SMI-S provider, check the SMI-S provider configuration. If the device was accidentally removed from the configuration, add it back, or if the SMI-S provider is not currently operational, restart the SMI-S provider. Then run a new IBM Spectrum Control SMI-S provider discovery job to allow IBM Spectrum Control to rediscover the SMI-S provider-to-device association. The existing IBM Spectrum Control performance monitor for the device should start working again automatically, once the SMI-S provider-to-device association is added back to the IBM Spectrum Control database by the discovery job.

HWNPM4091E Encountered an error during execution of a discover service process.

Explanation

An unexpected error occurred during discover process execution for the fiber channel switch.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4092E Encountered *exception* during execution of a discover service process.

Explanation

An unexpected error occurred during discover process execution for the fiber channel switch. The specified exception was caught as a result.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4093E An input business object could not be converted to a CIMInstance.

Explanation

An input business object (such as a Switch or a Port object) could not be converted to the CIMInstance form (required by an executing discover service process) because its CIM Keys were not found in the IBM Spectrum Control database. This could be because the input object was erroneous, hence no record for its keys exists in the database, or because the top level discovery is incomplete or erroneous (for example, the CIM keys of the Switch exist in the database, but not the keys of its constituent ports). This is an internal error.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4100E Failed to initialize SVC counter data service discover service reference.

Explanation

The SVC counter data service was unable to access the discover service. Without the discover service, the SVC counter data service cannot run.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4101E Failed to initialize SVC counter data service configuration service reference.

Explanation

The SVC counter data service was unable to access the configuration data service. Without the configuration data service, the SVC counter data service cannot run.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4102E Failed to parse performance data file time stamp suffix: *filename*.

Explanation

An error occurred while attempting to parse the time stamp suffix of an SVC iostats log file.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4103E SMI-S provider operation timeout (*timeout value* seconds) expired.

Explanation

The SMI-S provider communication timeout expired while executing an SMI-S provider operation.

Action

Increase the SMI-S provider communication timeout. If the problem persists, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4104E Failed to retrieve SMI-S provider password for SVC counter data service access point: *access point*.

Explanation

An error occurred while trying to retrieve the password associated with an SVC counter data service access point. Without the password, the SVC counter data service cannot login to the SMI-S provider for performance data collection.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4105E Encountered an error when communicating with the device agent.

Explanation

An unexpected error occurred when attempting to communicate with the SMI-S provider. The SMI-S provider is either down, or is not fully operational. This condition can occur if the network path is blocked by a firewall.

Action

Please ensure that the device agent is fully functional. It may be necessary to reboot the device agent, if it is in a hung state. Also ensure that a network path exists between the IBM Spectrum Control server and the device or device agent, including any appropriate firewall pass-throughs.

HWNPM4106E Encountered invalid SVC component type: *component type*.

Explanation

An invalid SVC component type was encountered while processing SVC performance data. This is an internal error.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4107E Failed to create performance data object: *performance data object class*.

Explanation

An error occurred while trying to create an instance of the specified performance data object. This is an internal error.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4108E TimeZone property is not defined for SVC cluster: *cluster identifier*.

Explanation

The TimeZone property is not defined for the specified SVC cluster. A timezone must be available for successful performance data retrieval.

Action

Set the TimeZone property on the SVC cluster specified in the message. If you have trouble doing so, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4109E SVC cluster TimeZone property is set to unrecognized value: *timezone id and name*.

Explanation

PM does not recognize the value associated with the SVC TimeZone property. This is an internal error.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4110E StatisticsStatus property is not defined for SVC cluster: *cluster identifier*.

Explanation

The StatisticsStatus property is not defined for the specified SVC cluster. This property must be set for PM to determine whether or not performance data collection is active on a given SVC. The value of this property is updated when performance data collection is either turned on or off. If it is not set, a problem might exist with the SVC.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4111E Failed to retrieve *dump filename* dump from SVC node *node identifier* (return code = *return code*).

Explanation

An error occurred while trying to retrieve the specified dump from the specified SVC node. As a result, performance data collection could not complete successfully. A problem might exist with the SVC.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4112E IsConfigNode property is not defined for SVC node: *node identifier*.

Explanation

The IsConfigNode property is not defined for the specified SVC node. The IsConfigNode property must be available for each node in a cluster for successful performance data retrieval.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4113E Caught *exception* while processing SVC XML performance data.

Explanation

An unexpected error occurred while attempting to parse SVC XML performance data.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4114E SVC cluster *cluster identifier* has more than one configuration node.

Explanation

The specified SVC cluster has more than one configuration node. The SVC is misconfigured.

Action

Configure the specified SVC cluster such that it has only one configuration node and retry performance data collection. If this message continues to appear, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4115E SVC cluster *cluster identifier* does not have a configuration node.

Explanation

The specified SVC cluster does not have a configuration node. The SVC is misconfigured.

Action

Configure the specified SVC cluster such that it has exactly one configuration node and retry performance data collection. If this message continues to appear, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4116W Failed to associate SVC performance data from non-configuration node with SVC performance data from configuration node.

Explanation

Time stamp information is used to correlate SVC performance data from non-configuration nodes with SVC performance data from configuration nodes. An attempt to perform such correlation failed. The likely cause is a mismatch in node clocks.

Action

Make sure the clocks of the nodes in the SVC from which data is being collected are in sync. If this message continues to appear, contact your IBM support representative.

HWNPM4117W Encountered incomplete SVC performance data sample.

Explanation

A SVC performance data sample is considered incomplete if one of the nodes in the cluster does not have an Nm_stats file. The likely cause of an incomplete performance data sample is a mismatch in node clocks.

Action

Make sure the clocks of the nodes in the SVC from which data is being collected are in sync. If this message continues to appear, contact your IBM support representative.

HWNPM4118E Firmware version information is not available for storage subsystem *subsystem name*. Performance data collection cannot proceed.

Explanation

Version information is not available for the specified storage subsystem.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4119E The firmware installed on storage subsystem *subsystem name (firmware version)* is not supported for performance data collection. The minimum level of firmware supported for performance data collection is *firmware version*.

Explanation

The firmware installed on the specified device is too old to perform performance monitoring.

Action

Follow the vendor's instructions to upgrade the firmware on the specified storage subsystem to the level specified in the message. If the problem persists after the upgrade, contact your IBM support representative.

HWNPM4150E Unable to retrieve storage subsystem for this device: *device identifier*.

Explanation

The storage subsystem was unable to be retrieved from the database.

Action

Check if the database contains the storage subsystem for this device.

HWNPM4151E Unable to determine the status of any performance data collection for this device: *device identifier*.

Explanation

An exception occurred while attempting to communicate with this device or its device agent.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4152E Performance data collection has already been enabled for this device: *device identifier*.

Explanation

Another user has already started performance data collection, and this device or its device agent only allows one collection to be performed at a time.

Action

Wait for the previously enabled collection to complete, then try again.

HWNPM4153E Performance data collection is not enabled for this device: *device identifier*.

Explanation

Performance data collection is not currently running for this device. It must be enabled before stopping it.

Action

None. Performance data collection is already stopped.

HWNPM4154E Unable to start performance data collection for this device: *device identifier*.

Explanation

An exception occurred while attempting to communicate with this device or its device agent.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4155E Failed to enable performance data collection for this device: *device identifier*.

Explanation

The device or device agent returned an error code when attempting to enable performance data collection.

Action

Check the frequency and duration of the performance data collection. If correct, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4156E Unable to stop performance data collection for this device: *device identifier*.

Explanation

An exception occurred while attempting to communicate with this device or its device agent.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4157E Failed to disable performance data collection for this device: *device identifier*.

Explanation

The device or device agent returned an error code when attempting to disable performance data collection.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4158E Unable to complete start performance data collection task for this device: *device identifier*.

Explanation

An exception occurred while attempting to execute an internal process.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4159E Unable to complete stop performance data collection task for this device: *device identifier*.

Explanation

An exception occurred while attempting to execute an internal process.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4160E Unable to complete performance data collection status query task for this device: *device identifier*.

Explanation

An exception occurred while attempting to execute an internal process.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4161E Performance data collection is not enabled for this device: *device identifier*.

Explanation

Performance data collection is not currently running for this device. It must be enabled in order to poll for performance data.

Action

Start performance data collection before polling for performance data.

HWNPM4162E Unable to retrieve port performance statistics data for this device: *device identifier*.

Explanation

The device or device agent returned an error code when attempting to retrieve port statistics.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4163E Unable to retrieve volume performance statistics data for this device: *device identifier*.

Explanation

The device or device agent returned an error code when attempting to retrieve volume statistics.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4164E Unable to retrieve rank performance statistics data for this device: *device identifier*.

Explanation

The device or device agent returned an error code when attempting to retrieve rank statistics.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4165E Unable to retrieve performance statistics data for this device: *device identifier*.

Explanation

An exception occurred while attempting to communicate with this device or its device agent.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4166E Unable to complete polling for performance data collection task for this device: *device identifier*.

Explanation

An exception occurred while attempting to execute an internal process.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4167E Unable to retrieve a device agent for this device: *device identifier*.

Explanation

An error occurred while attempting to obtain a device agent, such as an SMI-S provider, that controls this device.

Action

Confirm that this device is registered with a device agent, and that the device agent is known to IBM Spectrum Control. If there is a device agent, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4168E Failed attempt to use device *device identifier* counter data service with device *different device identifier*.

Explanation

A collection service is already assigned to another device.

Action

An internal error occurred with misassigned devices. Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4169E An invalid access point of device agent URL was used to acquire the agent for this device: *device identifier*.

Explanation

The access point passed into an internal process is in the wrong format.

Action

An internal error occurred with incorrect variables. Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4170E The device agent's configuration for *device identifier* has changed from the given access point, *device agent URL*.

Explanation

The access point passed into an internal process is incorrect.

Action

An internal error occurred with misassigned devices. Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4171E Performance data collection start task timed out after *time seconds* for device: *device identifier*.

Explanation

The internal task of starting performance data collection took too long.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4172E Performance data collection stop task timed out after *time seconds* for device: *device identifier*.

Explanation

The internal task of stopping performance data collection took too long.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4173E Performance data collection check status task timed out after *time* seconds for device: *device identifier*.

Explanation

The internal task of checking the performance data collection status took too long.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4174E Performance data collection poll task timed out after *time* seconds for device: *device identifier*.

Explanation

The internal task of polling for performance statistics data took too long.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4175W An error occurred while parsing statistics for port *port identifier*. Its statistics will be excluded.

Explanation

Unable to complete parsing statistics for this port. Possible problems include:

1. An internal component ID could not be created for the port.
2. A valid interval value was not returned.
3. A valid time stamp for the statistic was not returned.
4. The statistic returned was out of date.

Action

Contact your IBM support representative.

HWNPM4176W An error occurred while parsing statistics for volume *volume identifier*. Its statistics will be excluded.

Explanation

Unable to complete parsing statistics for this volume. Possible problems include:

1. An internal component ID could not be created for the volume.
2. A valid interval value was not returned.
3. A valid time stamp for the statistic was not returned.

4. The statistic returned was out of date.

Action

Contact your IBM support representative.

HWNPM4177W An error occurred while parsing statistics for rank *rank identifier*. Its statistics will be excluded.

Explanation

Unable to complete parsing statistics for this rank. Possible problems include:

1. An internal component ID could not be created for the rank.
2. A valid interval value was not returned.
3. A valid time stamp for the statistic was not returned.
4. The statistic returned was out of date.

Action

Contact your IBM support representative.

HWNPM4178E Failed to decrypt the device agent's password for device *device identifier*.

Explanation

An exception occurred while trying to decrypt the device agent's password. Without the password, the requested task cannot complete successfully.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4179W Performance data collection is currently enabled with errors for device *device identifier*.

Explanation

The device agent is indicating that the data collection has encountered a problem while collecting data from the device. This might result in the loss of some statistics data.

Action

Contact your IBM support representative.

HWNPM4180E Unable to retrieve *key identifier* value from the internal discover process.

Explanation

An internal process failed to map data correctly in order for the task to succeed.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4181W *number of ports of the port statistics from the device agent were unrecognized and were not included in this sample interval.*

Explanation

The port statistics data received from the device agent included data for a number of unknown ports. This could be due to one of the following conditions:

- The subsystem contains some ports which are not fibre-channel ports. IBM Spectrum Control as well as some device agents (mostly SMI-S providers) currently only support fibre-channel ports, so other types of ports (parallel SCSI-3, or ESCON ports for example) will be treated as unrecognized. If this is the case, this warning message can be safely ignored.
- The probe of the subsystem did not complete successfully, and failed to record information for all ports of the subsystem.
- The subsystem was recently upgraded and more ports were added which were not yet discovered by IBM Spectrum Control.

Action

If your subsystem contains unsupported ports, this message is expected but you can safely ignore it. Otherwise, run a probe for the subsystem to ensure that any unrecognized ports are discovered and properly recorded in the IBM Spectrum Control database. If this message persists for subsequent performance data collection intervals, even after successful completion of a probe, contact your IBM service representative.

HWNPM4182W *number of volumes of the volume statistics from the device agent were unrecognized and were not included in this sample interval.*

Explanation

The volume statistics data received from the device agent included data for a number of unknown volumes. This could be due to one of the following conditions:

- The probe of the subsystem did not complete successfully, and failed to record information for all volumes of the subsystem.
- The subsystem configuration was recently changed to add more volumes. These new volumes have not yet been discovered by IBM Spectrum Control.

Action

Rerun a probe for the subsystem to ensure that any unrecognized volumes are discovered and properly recorded in the IBM Spectrum Control database. If this message persists for subsequent performance data collection intervals, even after successful completion of a probe, contact your IBM service representative.

HWNPM4183W *number of ranks of the rank statistics from the device agent were unrecognized and were not included in this sample interval.*

Explanation

The rank statistics data received from the device agent included data for a number of unknown ranks. This could be due to one of the following conditions:

- The probe of the subsystem did not complete successfully, and failed to record information for all storage extents of the subsystem.
- The subsystem configuration was recently changed to add more ranks. These new storage extents have not yet been discovered by IBM Spectrum Control.

Action

Rerun a probe for the subsystem to ensure that any unrecognized ranks are discovered and properly recorded in the IBM Spectrum Control database. If this message persists for subsequent performance data collection intervals, even after successful completion of a probe, contact your IBM service representative.

HWNPM4184E *The device agent configured for this storage subsystem is not supported for this task. The current version, version number, is downlevel from from the minimum required, version number.*

Explanation

Performance data cannot be collected from the subsystem through this particular device agent (usually an SMI-S provider), because the version of the agent is not supported by the IBM Spectrum Control Performance Manager.

Action

Use another device agent, or upgrade the version of the existing device agent to the minimum required level indicated in the message.

HWNPM4185W The device agent did not return all performance statistics data for this time interval. The incomplete data is being processed.

Explanation

IBM Spectrum Control expects to receive performance data for ports, ranks, and volumes from the device and device agent. One or two of these types of data was not received as expected. This usually indicates that there is a problem with the device itself, the native device interfaces, or the device agent.

In the rare case that there are no ranks or no volumes defined on the device, this message can safely be ignored because then no rank or volume statistics will be sent by the device. However, in that case it is not necessary to run a performance monitor for the device because there is no performance to be measured.

Action

Ensure that the device and the native device interfaces are operating normally. Ensure that the device agent is operating normally. It might be necessary to view error and/or trace logs for your device and/or device agent. Contact your IBM service representative for help, if necessary.

HWNPM4186W The ESS SMI-S provider did not return performance statistics data for both clusters for this time interval. The incomplete data is being processed.

Explanation

IBM Spectrum Control expects to receive performance data for ports, ranks, and volumes from ESS SMI-S provider for both clusters. The data for one cluster was not received as expected. This usually indicates that there is a problem with the device itself, the native device interfaces, or the device agent.

Action

Ensure that the device and the native device interfaces are operating normally. Ensure that the ESS SMI-S provider is operating normally. It might be necessary to view error and/or trace logs for your device and/or device agent. Contact your IBM service representative for help, if necessary.

HWNPM4187W The device does not support performance management for pool *pool ID* because it contains Space Efficient Volumes. Only incomplete performance data can be collected for array *array ID*.

Explanation

The specified pool contains space efficient volumes, which makes it impossible to accurately manage the performance for those ranks, the arrays associated with those ranks, and the device adapters associated with those arrays.

For DS8000 devices whenever a pool consisting of multiple ranks contains space efficient volumes that are not yet fully allocated, the performance impact of those volumes on their associated ranks cannot be measured. As a result, to avoid presenting the user with potentially inaccurate or misleading performance data, the Performance Manager does not attempt to compute the performance metrics for the affected arrays and device adapters

Action

None.

HWNPM4188W The performance monitor was unable to collect performance statistics data from the device agent for the following component types: *component list*.

Explanation

The request for performance statistics from the agent resulted in no statistic data being received for either the ports, ranks, and/or volumes.

The performance monitor will attempt to retrieve statistics for all components until the next polling interval. If the components statistics continue to be missing for the next polling interval, the statistics that are available will be processed. This may result in the loss of some performance data.

Action

Confirm that both the subsystem and the subsystem's agent are working properly. Contact your IBM support representative for further assistance.

HWNPM4189W *number of MDisks of the MDisk statistics from the device agent were unrecognized and were not included in this data collection interval.*

Explanation

The MDisk statistics data received from the device agent included data for a number of unknown MDisks. This could be due to one of the following conditions:

- The probe of the subsystem did not complete successfully, and failed to record information for all storage extents of the subsystem.
- The subsystem configuration was recently changed to add more MDisks. These new storage extents have not yet been discovered by IBM Spectrum Control.

Action

Rerun a probe for the subsystem to ensure that any unrecognized MDisks are discovered and properly recorded in the IBM Spectrum Control database. If this message persists for subsequent performance data collection intervals, even after successful completion of a probe, contact your IBM service representative.

Related reference

-  [Getting support](#)

HWNPM4190W *number of nodes of the node statistics from the resource agent were unrecognized and were not included in this data collection interval.*

Explanation

The node statistics data that were received from the data source included data for a number of unknown nodes. This problem might be caused by the following conditions:

- The probe of the storage system did not complete successfully, and failed to record information for all nodes of the storage system.
- The storage system configuration was recently changed to add more nodes. These new nodes were not discovered.
- If the storage system is an IBM Spectrum Scale cluster, the OS host names and the cluster host names of the Spectrum Scale nodes might conflict. Similarly the cluster host names and the Zimon sensor or collector configuration files might conflict.

Action

If the storage system is an IBM Spectrum Scale cluster, check the OS host name and the IBM Spectrum Scale cluster host name of each of the nodes. Correct any mismatches, for example use the `hostnamectl` command to change the OS host name to the one shown by `mmlscluster`, which is the host name that is recognized by the cluster.

If the storage system is an IBM Spectrum Scale cluster, check the Zimon sensor configuration files on each of the nodes and the Zimon collector configuration file on the collector node. Correct any mismatches in the node host names that are used in those configuration files.

Otherwise, rerun a probe for the storage system to ensure that any unrecognized nodes are discovered and properly recorded in the database repository. If this message persists for subsequent performance data collection intervals, even after successful completion of a probe, contact IBM Software Support.

HWNPM4191W *number of modules out of total number of modules module statistics could not be retrieved from the device agent due to errors, and were not included in this data collection interval.*

Explanation

Each interface module in an XIV device is responsible for tracking performance statistics of the I/O flowing through it. In this case, one or more of the interface modules did not return its statistics data to IBM Spectrum Control. This is usually caused by the particular modules being inoperative.

Action

Check the health status of your XIV interface modules, and correct any errors. If this message persists for subsequent performance data collection intervals, even if all XIV interface modules are fully operational, contact your IBM service representative.

HWNPM4192W *number of Drives of the drive statistics from the device agent were unrecognized and were not included in this data collection interval.*

Explanation

The performance data gathered for local disks or flash modules from the data source included data for a number of unknown drives.

This warning can appear if one of the following conditions occurs:

- The probe of the storage system did not complete successfully, and failed to record information for all local disks or flash modules of the storage system.
- The storage system configuration was recently changed to add more drives. The new drives were not yet discovered when the previous probe was run.

Action

Run another probe for the storage system to ensure that any unrecognized local disks and flash modules are discovered and properly recorded in the database repository. If this message continues after a probe completes successfully, contact IBM Software Support.

HWNPM4193W *number of Volume-copies of the volume-copy statistics from the device agent were unrecognized and were not included in this data collection interval.*

Explanation

The volume-copy statistics data received from the device agent included data for a number of unknown copies. This could be due to one of the following conditions:

- The probe of the subsystem did not complete successfully, and failed to record information for all volume-copies of the subsystem.
- The subsystem configuration was recently changed to add more volumes or volume-copies. These new volumes or volume-copies have not yet been discovered by IBM Spectrum Control.

Action

Rerun a probe for the subsystem to ensure that any unrecognized volume-copies are discovered and properly recorded in the IBM Spectrum Control database. If this message persists for subsequent performance data collection intervals, even after successful completion of a probe, contact your IBM service representative.

Related reference

-  [Getting support](#)

HWNPM4194W *number of partitions of the partition statistics from the device agent were unrecognized and were not included in this data collection interval.*

Explanation

The partition statistics data received from the device agent included data for a number of unknown pools. This could be due to one of the following conditions:

- The probe of the subsystem did not complete successfully, and failed to record information for all pools of the subsystem.
- The subsystem configuration was recently changed to add more pools. These new pools have not yet been discovered by IBM Spectrum Control.

Action

Rerun a probe for the subsystem to ensure that any unrecognized pools are discovered and properly recorded in the IBM Spectrum Control database. If this message persists for subsequent performance data collection intervals, even after successful completion of a probe, contact your IBM service representative.

Related reference

-  [Getting support](#)

HWNPM4195W *number of file systems* of the file system statistics from the device agent were unrecognized and were not included in this data collection interval.

Explanation

The file system statistics data received from the device agent included data for a number of unknown file systems. This could be due to one of the following conditions:

- The probe of the storage system did not complete successfully, and failed to record information for all file systems of the storage system.
- The storage system configuration was recently changed to add more file systems. These new file systems have not yet been discovered by IBM Spectrum Control.

Action

Rerun a probe for the storage system to ensure that any unrecognized file systems are discovered and properly recorded in the IBM Spectrum Control database. If this message persists for subsequent performance data collection intervals, even after successful completion of a probe, contact your IBM service representative.

Related reference

- [Getting support](#)

HWNPM4250E Failed to start the discover service for the SMI-S counter data service.

Explanation

An error occurred while attempting to start the discover service for the SMI-S counter data service.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4251E Failed to start the configuration service for the SMI-S counter data service.

Explanation

An error occurred while attempting to start the configuration service for the SMI-S counter data service.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4252I Successfully returned access point device *namer* for device *device name*.

Explanation

The access point required for the SMI-S counter data service was successfully retrieved.

Action

None.

HWNPM4253I Successfully stopped SMI-S counter data service on access point *access point* for device *device name*.

Explanation

The SMI-S counter data service was successfully stopped.

Action

None.

HWNPM4254I The SMI-S counter data service is active on access point *access point* for device *device name*.

Explanation

The SMI-S counter data service is active.

Action

None.

HWNPM4255I The SMI-S counter data service is inactive on access point *access point* for device *device name*.

Explanation

The SMI-S counter data service is inactive.

Action

None.

HWNPM4256I Performance statistics successfully returned on access point *access point* for device *device name*.

Explanation

The performance statistics for the SMI-S counter data service were successfully returned.

Action

None.

HWNPM4257W Performance statistics not returned on access point *access point* for device *device name*.

Explanation

The performance statistics for the SMI-S counter data service were not successfully returned.

Action

If SMI-S block storage performance statistics are supported on the device, then contact your IBM support representative.

HWNPM4258E No SMI-S providers found for device *device name*.

Explanation

No SMI-S providers were found for the device. The SMI-S provider is required to retrieve performance statistics for the SMI-S counter data service.

Action

Check if an SMI-S provider is defined for the device. If the SMI-S provider is defined, then contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4259E No storage subsystem found for device *device name*.

Explanation

No storage subsystem was found for the device. A storage subsystem is required to retrieve performance statistics for the SMI-S counter data service.

Action

Check if a storage subsystem is defined for the device. If the storage subsystem is defined, then contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4260E Failed to initialize the polling context for device *device name*.

Explanation

The polling context contains the data required to retrieve storage subsystem statistics from the SMI-S counter data service. Therefore, the request fails.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4261E Failed to retrieve the device capabilities for device *device name*.

Explanation

The attempt to retrieve the device capabilities for the device failed.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4262E A database exception occurred trying to retrieve the device capabilities for device *device name*.

Explanation

Unable to retrieve the device capabilities for the device due to a database exception.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4263E A database exception occurred trying to retrieve the storage subsystem for device *device name*.

Explanation

Unable to retrieve the storage subsystem for the device due to a database exception.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4264W Failed to retrieve manifest for *component type*.

Explanation

The attempt to retrieve a manifest for the device failed. Each component type will have its own manifest, and if the component is not supported, the manifest will not be located.

Action

If the component type is supported, contact your IBM support representative.

HWNPM4265E A database exception occurred trying to retrieve the Manifests for device *device name*.

Explanation

Unable to retrieve the manifests for the device due to a database exception.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4266E No manifests found for device *device name*.

Explanation

No manifests were found for the device. The manifests are required to parse the performance statistics that are returned. Therefore, the request fails.

Action

If the component type is supported, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4267E A database exception occurred trying to retrieve the discovery parameters for device *device name*.

Explanation

Unable to retrieve the discovery parameters the device due to a database exception.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4268E Statistics record not correctly formatted due to exception *local exception string*.

Explanation

The statistics record is not correctly formatted.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4269E Statistics record not correctly parsed due to exception *local exception string*.

Explanation

The statistics record is not correctly parsed.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4270W The block storage statistics is not formatted for device *device name*.

Explanation

The block storage statistics is not correctly formatted for the device.

Action

If block storage statistics are supported for this device, contact your IBM support representative.

HWNPM4271E The SMI-S provider found for device *device name* is not valid.

Explanation

The SMI-S provider found for the device is not valid. The SMI-S provider is required to retrieve performance statistics for the SMI-S counter data service.

Action

Check if an SMI-S provider found for the device is valid. If the SMI-S provider is valid, then contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4272E The storage subsystem found for device *device name* is not valid.

Explanation

The storage subsystem found for the device is not valid. The storage subsystem must have a serial number, and a storage subsystem is required to retrieve performance statistics for the SMI-S counter data service.

Action

Check if a storage subsystem has a defined serial number. If the storage subsystem has a defined serial number, then contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4273W Discarding the stale performance statistics returned on access point *access point* for device *device name*.

Explanation

Some of the performance statistics counters were not updated on the SMI-S provider for the device. IBM Spectrum Control will discard all the performance data retrieved from the SMI-S provider during this sample interval. IBM Spectrum Control will automatically try to retrieve the performance statistics after a short timeout.

Action

None

HWNPM4274E The SMI-S provider found for this device has changed. Please re-run SMI-S provider discovery and probe.

Explanation

The SMI-S provider may have been migrated from the Brocade SMI Agent to the SMI Agent that is integrated in Brocade Data Center Fabric Manager, or vice versa.

Action

If SMI-S provider discovery and probe are executed after SMI-S provider migration and the problem still exists, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM4300E Access to the agent or device has been denied. Ensure that valid credentials have been specified for agent *agent name*.

Explanation

When attempting to retrieve performance statistics for the device, access was denied by the device or the device agent. Therefore no performance statistics could be retrieved. No performance data will be inserted for the device in this time period. The next performance data sample recorded into the database might represent an average over more than the configured interval length.

Action

Ensure that the correct agent address has been specified, and that the specified credentials are valid and will allow access to the device or agent for performance data collection. The credentials usually consist of a username and password, but can also encompass other security related parameters such as ssh keys or authentication tokens, depending on the type of device or agent being accessed and, where applicable, the access method selected.

HWNPM4301E The device or device agent did not respond within the allotted time (*timeout value seconds*).

Explanation

When attempting to retrieve performance statistics for the device, the requested performance data was not returned before the timeout expired. Either the device or agent is completely unresponsive, or is much slower than expected by IBM Spectrum Control. No performance data will be inserted for the device in this time period. The next performance data sample recorded into the database might represent an average over more than the configured interval length.

Action

Ensure that the device and device agent are fully operational. It may be necessary to reboot either device or agent, if it is in a hung state. If you have reason to believe that the device and agent are operational but are simply slower than expected, you can also attempt to increase the timeout value used by the performance manager, which is set in configuration file `device/conf/pm.conf`.

HWNPM4302E New performance data is not yet available for the device. Statistics with time stamps later than *time_stamp* could not be found.

Explanation

If the time stamp in the message is "null", then no statistics were previously retrieved, and the performance manager is unable to get ANY statistics for the device.

In case performance data is cached by the device or device agent, the performance manager ensures that the most recently retrieved performance statistics are indeed newer than the previously retrieved statistics. If not, the performance manager waits for a short time and tries retrieving the statistics again. If after several tries no new statistics were retrieved, this message is issued.

No performance data is inserted for the device in this time period. The next performance data sample that is recorded into the database might represent an average over more than the configured interval length.

Action

Ensure that the device and device agent are fully operational. It might be necessary to restart either device or agent, if it is in a hung state. Ensure that if the device has multiple clocks (for example for multiple nodes or controllers), that the clocks are synchronized to within a few minutes.

HWNPM4303E An agent API call (*API name*) failed while attempting to retrieve performance data for the device.

Explanation

The current attempt to retrieve a set of performance data from the device failed. No performance data will be inserted for the device in this time period. The immediate operation of the running performance monitor is unaffected. However the next performance data sample recorded into the database might represent an average over more than the configured interval length.

The API name indicated in the message is the function or method which caused the error. If a null name is indicated, the exact function or method is unknown.

Action

Ensure that the device agent is fully operational. It may be necessary to reboot the agent if it is not working properly. If the problem persists, contact the vendor of the device and agent for problem diagnosis and resolution.

HWNPM4304E The request for performance data could not be retrieved from the queue by the data collector.

Explanation

The data collector may be down, or the connection may be unavailable or congested.

Action

Verify the connection between the data collector and the server. Then, verify that the data collector is running and if not, restart the data collector process.

HWNPM4305W No samples were received from the data collector in the expected time. The data might still arrive automatically after connection is recovered.

Explanation

The data collector might be down, or the connection might not be available or congested. The data might still arrive automatically after connection is recovered.

Action

Verify the connection between the data collector and the server. Then, verify that the data collector is running and if not, restart the data collector process. If the data collector is running, restart data collection. The data might still arrive automatically after connection is recovered.

HWNPM4306E The data collector failed to connect to the storage management service because of invalid credentials. No performance manager data can be collected from *device name* until valid credentials are available.

Explanation

The credentials of the data collector were not authenticated by the storage management service. No further automated attempts to connect to the storage management service can be made until the valid credentials are available.

Action

Update the credentials and restart the data collector. To restart the data collector, log in to the computer that the data collector and its file are on. On a UNIX system, run the dataCollector.sh script to start the data collector. On a Windows system, start the service from the Services tab on the Windows Task Manager. For more information, see Resolving Connections Issues topic in the Knowledge Center.

HWNPM4502E Attempt to delete a default policy.

Explanation

An IBM default policy or current default policy cannot be deleted.

Action

None.

HWNPM4503E A database update operation failed.

Explanation

A failure occurred when IBM Spectrum Control attempted to make an update to the database.

Action

Verify that database is operational and online.

HWNPM4504E A database insert operation failed.

Explanation

A failure occurred when IBM Spectrum Control attempted to insert a new row into the database.

Action

Verify that database is operational and online.

HWNPM4505E A database delete operation failed.

Explanation

A failure occurred when IBM Spectrum Control attempted to delete one or more rows from a database table.

Action

Verify that database is operational and online.

HWNPM4506E A database cursor operation failed.

Explanation

A failure occurred when IBM Spectrum Control attempted to obtain or use a database cursor.

Action

Verify that database is operational and online.

HWNPM4507E A database connect operation failed.

Explanation

A failure occurred when IBM Spectrum Control attempted to connect to the database.

Action

Verify that database is operational and online.

HWNPM4508E A database retrieve operation failed.

Explanation

A failure occurred when IBM Spectrum Control attempted to retrieve a row from a database table.

Action

Verify that database is operational and online.

HWNPM4509E A database operation failed.

Explanation

A generalized database failure occurred.

Action

Verify that database is operational and online.

HWNPM4510E A database query operation failed.

Explanation

A failure occurred when IBM Spectrum Control attempted a database query operation.

Action

Verify that database is operational and online.

HWNPM4511E A database commit operation failed.

Explanation

A failure occurred when IBM Spectrum Control attempted a database commit operation.

Action

Verify that database is operational and online.

HWNPM5200E The performance manager failed to publish event even name due to exception exception.

Explanation

The performance manager failed to publish event to other modules.

Action

Restart device server and try again.

HWNPM5210E The performance manager failed to receive event from other modules.

Explanation

The performance manager failed to receive event from other modules.

Action

Restart device server and try again.

HWNPM5211E The first parameter passed to this method is null.

Explanation

The subscriber can not be null when calling this method.

Action

Make sure the first parameter passed to this method is not null.

HWNPM5212E The second parameter passed to this method is invalid.

Explanation

The event array can not be null or empty.

Action

Make sure the second parameter passed to this method is valid.

HWNPM5400E The performance data collection identifiers are not valid integers: schedule ID {0}, schedule run number {1}, job run number {2}.

Explanation

Each performance data collection job is identified by three integers. The scheduler service did not receive three integers when asked to update the job status, and so could not identify the correct job.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM5401E There was a problem establishing the database connection: {0}.

Explanation

An exception occurred when trying to retrieve the database connection from the connection pool.

Action

Make sure your database is working correctly. If you cannot find anything wrong with the database, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM5402E There was a problem creating the new run job entry: {0}.

Explanation

An exception occurred when trying to insert a new run job entry into the database.

Action

Make sure your database is working correctly. If you cannot find anything wrong with the database, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM5403E There was a problem updating the run job entry {0}: {1}.

Explanation

An exception occurred when trying to update an existing run job entry in the database.

Action

Make sure your database is working correctly. If you cannot find anything wrong with the database, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM5404E There was a problem closing the database connection: {0}.

Explanation

An exception occurred when trying to close the database connection.

Action

Make sure your database is working correctly. If you cannot find anything wrong with the database, contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM5405E There was a problem inserting a new run job into the database: {0}.

Explanation

An exception occurred when trying to insert the run job into the database table T_RUN_JOBS.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM5406E There was a problem executing an update for run job number {0} in the database.

Explanation

No rows were updated when an attempt was made to update a run job in the database.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM5407E There was a problem executing an update for run job number {0} in the database.

Explanation

No rows were updated when an attempt was made to update a run job in the database.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM5408E There was a problem executing an update for run number {0} in the database.

Explanation

No rows were updated when an attempt was made to update a run in the database.

Action

Contact your IBM support representative.

Related reference

- [Getting support](#)

HWNPM5409I Successfully retrieved the configuration data for the elastic device. Found *number of nodes* Nodes and *number of file systems* File systems,

Explanation

The running performance monitor has updated its internal state with the latest configuration data from the storage subsystem. The indicated device components were found.

Action

None.

HWNPM5410W The performance monitor could not collect performance data for the following cluster nodes: *nodes names*.

Explanation

The IBM Spectrum Scale performance monitoring tool on the GPFS cluster could not collect performance data for the nodes and related resources such as file systems.

Action

Check that the collector component of the performance monitoring tool is started on a single GPFS cluster node.

Check that the sensor component of the performance monitoring tool is started on the nodes.

Check the sensor and collector components are configured correctly on the nodes. For information about how to configure the performance monitoring tool, see <https://www.ibm.com/docs/en/spectrum-scale/5.1.4?topic=tool-configuring-performance-monitoring>.

Try the operation again.

HWNPM5411W The performance monitor could not collect performance data for the following filesystems: *filesystem names*.

Explanation

The IBM Spectrum Scale performance monitoring tool on the GPFS cluster could not collect performance data for the file systems. This problem might occur because the performance monitoring tool is not started or is not configured correctly on the nodes that the file system is mounted on.

Action

Check that the collector component of the performance monitoring tool is started on a single GPFS cluster node.

Check that the sensor component of the performance monitoring tool is started on the nodes that the file system is mounted on.

Check the sensor and collector components are configured correctly on the nodes. To learn more about configuring the performance monitoring tool, see http://www.ibm.com/support/knowledgecenter/STXKQY_4.1.1/com.ibm.spectrum.scale.v4r11.adv.doc/bl1adv_PMToverview.htm

Try the operation again.

HWNPM5412E Performance statistics collection is not enabled.

Explanation

Performance statistics collection is not enabled on either the SMI-S provider or the system it manages.

Action

Enable performance statistics collection on either the SMI-S provider or the system it manages.

HWNPM5413E The process failed because the userid or password provided failed to connect to the Export Tool.

Explanation

Performance data collection for Hitachi systems requires a user name and password to be defined in the Hitachi Device Manager for the Export Tool. Connection with the provided user name or password failed.

Action

Check the credentials for the user that is configured to monitor the block storage system. To connect to the storage system for both probe and performance monitoring, ensure that the following credentials match:

- The user name and password that is defined in Hitachi Device Manager.
- The user name and password that is used for Hitachi Command Suite to connect to the device.

HWNPM5414E The process failed because the Hitachi SVP was busy and did not return data or timed out.

Explanation

The Hitachi external process was unable to collect data because the Hitachi SVP is too busy.

Action

The Hitachi SVP seems to be too busy to return performance data. Collection will retry at the next performance interval.

HWNPM5415E The process failed because the performance interval is set to 0.

Explanation

Incorrect data can cause IBM Spectrum Control to set the performance interval to 0.

Action

Collect service logs for the product. IBM Software Support can use these logs to help troubleshoot the issue. Open a support case at <https://www.ibm.com/mysupport/> for help with determining what might be causing this error. Don't forget to include the version, release, modification, and service level number of IBM Spectrum Control that you're using.

For information about how to collect service logs, go to the documentation at <https://www.ibm.com/docs/en/spectrum-control/5.4.4?topic=overview-collecting-service-logs-software-support-troubleshooting>.

HWNPM5416E The process failed because the performance interval for the storage system is not supported.

Explanation

For Hitachi Virtual Storage Platform (VSP) F and G Series storage systems, you can select performance intervals of 1 minute and 5 minutes. Other performance intervals are not supported.

Action

For the storage system, set the performance interval to 1 minute or 5 minutes. To set the interval, go to Storage > Block Storage Systems, right-click the storage system, and select Data Collection > Schedule.

If performance monitoring is already enabled at an unsupported interval and you want to keep the previous monitoring data, first export the data by using the Hitachi Export Tool and then restart monitoring.

HWNPM5417E The process failed because the Hitachi VSP Model being monitored is not known. A Hitachi Export Tool to match it cannot be found.

Explanation

This problem might occur when the appropriate version of the Hitachi Export Tool is not installed for the model of the storage system. The Hitachi Export Tool is used to collect performance metadata about the storage system and is located on the server or virtual machine where the data collector is installed.

Action

Ensure that the configured Export Tool matches the model of the Hitachi VSP storage system. If the appropriate Export Tool is not being used, log in with your Hitachi account to download the correct version from the following location:
https://knowledge.hitachivantara.com/Knowledge/Storage/How_to_Download_the_Appropriate_Export_Tool_Version_Specific_to_Array_Microcode.

For information about installing the Hitachi Export Tool, go to the documentation at <https://www.ibm.com/docs/en/storage-insights?topic=pro-installing-hitachi-export-tool>.

HWNPM5418E The process failed because the data collected is out of range.

Explanation

Performance data that is returned from the device is out of range and not valid.

Action

If the problem persists, open a support case at <https://www.ibm.com/mysupport/>. Don't forget to include the version, release, modification, and service level number of IBM Spectrum Control that you're using.

HWNPM5419E Performance data can't be collected.

Explanation

No other information is available.

Action

If this problem persists, open a support case at <https://www.ibm.com/mysupport/>. Don't forget to include the version, release, modification, and service level number of IBM Spectrum Control that you're using.

HWNRM - Replication manager messages

- [HWNRM0000I Connection to Replication Manager Server successful.](#)
- [HWNRM0001E Communication with Replication Manager server failed.](#)
- [HWNRM0002E The specified port number {0} is invalid.](#)
- [HWNRM0003E Invalid host or port specified.](#)
- [HWNRM0004E Unknown host error.](#)
- [HWNRM0005E Failed to connect to the Replication Manager server.](#)
- [HWNRM0006E Status update failed in database.](#)
- [HWNRM0007E Unable to read status from database.](#)
- [HWNRM0008E Certificate file for authentication with Replication Manager not found.](#)
- [HWNRM0009E Replication Manager certificate error.](#)
- [HWNRM0010E Unable to read Replication Manager authentication certificate.](#)
- [HWNRM0100E The delete action failed because the session check with the Replication Manager server failed. {0}](#)
- [HWNRM0101E One or more storage systems cannot be deleted because they contain volumes that are defined in an active replication session. {0}](#)
- [HWNRM0102W The volume is in active replication session. {0}](#)
- [HWNRM0103W Error getting storage subsystems.](#)
- [HWNRM0104E Error getting storage subsystem information.](#)
- [HWNRM0105E Storage subsystem not found in IBM Spectrum Control database.](#)
- [HWNRM0106E Storage subsystem type is not supported for Replication.](#)
- [HWNRM0107E Volume not found in the IBM Spectrum Control database.](#)
- [HWNRM0108E Volume not found, volume not valid. {0}](#)
- [HWNRM0109E Resource not found in Replication Manager. {0}](#)
- [HWNRM0110E Error getting volume information.](#)
- [HWNRM0200E Error getting server information from the database.](#)
- [HWNRM0201E Error updating the server information into the database.](#)
- [HWNRM0011E Replication Manager server is not installed.](#)
- [HWNRM0012E Adding the connection to the specified storage device failed on the Replication server](#)
- [HWNRM0013E The connection specified to be added to the Replication server is not a valid one. Please check the parameters again](#)
- [HWNRM0014E A problem with the following message Exception message appeared when modifying the connection on the Replication server](#)
- [HWNRM0015E The connection to be modified does not exist on the Replication server](#)

HWNRM0000I Connection to Replication Manager Server successful.

Explanation

IBM Spectrum Control can successfully communicate with the Replication Manager server.

Action

None.

HWNRM0001E Communication with Replication Manager server failed.

Explanation

IBM Spectrum Control can not communicate with the Replication Manager server.

Action

Verify correct port is used for Replication Manager server and that the Replication Manager server is up and running.

HWNRM0002E The specified port number {0} is invalid.

Explanation

The port number specified is not correct.

Action

Use numeric port number.

HWNRM0003E Invalid host or port specified.

Explanation

The hostname and/or the port number for the Replication Manager server is invalid.

Action

Provide a valid hostname/port number for the Replication Manager server. If only port number can be changed, provide valid port number.

HWNRM0004E Unknown host error.

Explanation

Can not resolve host for the Replication Manager server.

Action

If hostname doesn't work, try providing the IP address.

HWNRM0005E Failed to connect to the Replication Manager server.

Explanation

The most likely cause of this error is that wrong port number is given for the Replication Manager Server. Another possible cause of this error is the internal authentication failure between IBM Spectrum Control and the Replication Manager Server.

Action

Make sure that the port number entered is correct. The correct port number can be found from the config file repcli.properties located at {replication manager install directory}/replication/CLI/. The default port is 5110. Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0006E Status update failed in database.

Explanation

Could not update the status of Replication Manager connection test in the IBM Spectrum Control database.

Action

Make sure IBM Spectrum Control is able to communicate with the database. If the problem persists, contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0007E Unable to read status from database.

Explanation

Could not read the previous status of Replication Manager connection test in the IBM Spectrum Control database.

Action

Make sure IBM Spectrum Control is able to communicate with the database, if problem persists, contact IBM service.

Related reference

- [Getting support](#)

HWNRM0008E Certificate file for authentication with Replication Manager not found.

Explanation

IBM Spectrum Control could not locate the certificate file to connect to the Replication Manager server.

Action

This could be due to failed installation of Replication Manager server, or deleting some files from the IBM Spectrum Control installation directory.

HWNRM0009E Replication Manager certificate error.

Explanation

IBM Spectrum Control found an invalid/corrupted certificate file which is needed to communicate with the Replication Manager server.

Action

Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0010E Unable to read Replication Manager authentication certificate.

Explanation

IBM Spectrum Control found an invalid/corrupted certificate file which is needed to communicate with the Replication Manager server.

Action

Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0100E The delete action failed because the session check with the Replication Manager server failed. {0}

Explanation

A storage subsystem or volume can not be deleted without successfully checking if it is in an active Replication session.

Action

Make sure Replication Server is up and running and do a connectivity test to it from the IBM Spectrum Control.

Related reference

- [Getting support](#)

HWNRM0101E One or more storage systems cannot be deleted because they contain volumes that are defined in an active replication session. {0}

Explanation

You cannot delete storage systems if they contain volumes that are defined in active replication sessions. In IBM Spectrum Control for Replication, a replication session is used to perform a specific type of data replication against a specific set of volumes.

Action

Go to Replication Manager and remove all the volumes of the storage systems from the associated replication sessions. To access Replication Manager from the IBM Spectrum Control GUI, go to Replication Manager > Replication Management in the navigation tree and click Replication Sessions Overview. When the volumes are removed from the active replication sessions, try to delete the storage systems again.

HWNRM0102W The volume is in active replication session. {0}

Explanation

A storage subsystem volume can not be deleted if it is in an active Replication session.

Action

Go to Replication Manager and remove the storage subsystem volume from the Replication session.

HWNRM0103W Error getting storage subsystems.

Explanation

Not able to retrieve storage subsystems from Replication Manager server.

Action

Make sure Replication Manager server is up and running and that the IBM Spectrum Control is able to communicate with the Replication Manager server.

HWNRM0104E Error getting storage subsystem information.

Explanation

Not able to retrieve storage subsystem details from Replication Manager server.

Action

Make sure Replication Manager server is up and running and that IBM Spectrum Control is able to communicate with the Replication Manager server.

HWNRM0105E Storage subsystem not found in IBM Spectrum Control database.

Explanation

Not able to retrieve storage subsystem details from the IBM Spectrum Control database.

Action

Make sure IBM Spectrum Control is able to communicate with the database. If problem persists, contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0106E Storage subsystem type is not supported for Replication.

Explanation

The given storage subsystem type is not supported by Replication Manager.

Action

Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0107E Volume not found in the IBM Spectrum Control database.

Explanation

Volume not found in the IBM Spectrum Control database.

Action

The specified volume not found in the IBM Spectrum Control database. If problem persists, please contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0108E Volume not found, volume not valid. {0}

Explanation

Volume not found, volume not valid.

Action

The specified volume not found in the IBM Spectrum Control database. If problem persists, please contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0109E Resource not found in Replication Manager. {0}

Explanation

The storage subsystem or volume not found in Replication Manager.

Action

Contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0110E Error getting volume information.

Explanation

Error getting volume information from the IBM Spectrum Control database.

Action

Make sure IBM Spectrum Control is able to connect to database. If problem persists, contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0200E Error getting server information from the database.

Explanation

Error getting Replication Manager server information from the IBM Spectrum Control database.

Action

Make sure IBM Spectrum Control is able to connect to database. If problem persists, contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0201E Error updating the server information into the database.

Explanation

Error updating Replication Manager server information into the IBM Spectrum Control database.

Action

Make sure IBM Spectrum Control is able to connect to database. If problem persists, contact IBM customer technical support with all related errors.

Related reference

- [Getting support](#)

HWNRM0011E Replication Manager server is not installed.

Explanation

IBM Spectrum Control has not detected Replication Manager server installation.

Action

If Replication Manager functionality is required, make sure it is installed and running. Replication Manager server should be installed using the IBM Spectrum Control installer.

HWNRM0012E Adding the connection to the specified storage device failed on the Replication server

Explanation

Adding the connection to the specified storage device failed

Action

HWNRM0013E The connection specified to be added to the Replication server is not a valid one. Please check the parameters again

Explanation

The connection specified is not a valid one

Action

HWNRM0014E A problem with the following message *Exception message* appeared when modifying the connection on the Replication server

Explanation

A problem appeared when modifying the connection

Action

HWNRM0015E The connection to be modified does not exist on the Replication server

Explanation

The connection to be modified does not exist on the Replication server

Action

HWNSS - Single sign-on User Interface messages

- [HWNSS0001E The IBM Spectrum Control device server is down and cannot perform OS user authentication. It is still possible to perform OS user authentication against the data server, however since the device server is down the IBM Spectrum Control functionality will be limited. Among the](#)

limitations is the inability to perform SSO to other applications that rely on the presence of a lightweight third party authentication token. To proceed enter a local OS username and password.

- HWNSS0002E The IBM Spectrum Control device server is down and cannot perform LDAP user authentication.
- HWNSS0003E The single sign-on (SSO) token is missing or incorrect. Enter a valid user name and password.
- HWNSS0004E The web server is unavailable and cannot be used for authentication. The Device server can be used for some authentication, but IBM Spectrum Control functions will be limited. To proceed, enter either the common user name that was used to install IBM Spectrum Control or the tpcFileRegistryUser user name and password.

HWNSS0001E The IBM Spectrum Control device server is down and cannot perform OS user authentication. It is still possible to perform OS user authentication against the data server, however since the device server is down the IBM Spectrum Control functionality will be limited. Among the limitations is the inability to perform SSO to other applications that rely on the presence of a lightweight third party authentication token. To proceed enter a local OS username and password.

Explanation

Authentication and Single Sign On (SSO) functionality is primarily obtained through the IBM Spectrum Control Device Server. With only the IBM Spectrum Control Data Server running, IBM Spectrum Control can only authenticate the user with no SSO capability.

Action

Restart the IBM Spectrum Control Device Server and re-login to IBM Spectrum Control.

HWNSS0002E The IBM Spectrum Control device server is down and cannot perform LDAP user authentication.

Explanation

LDAP authentication is only available through the IBM Spectrum Control Device Server.

Action

Restart the IBM Spectrum Control Device Server and re-login to IBM Spectrum Control.

HWNSS0003E The single sign-on (SSO) token is missing or incorrect. Enter a valid user name and password.

Explanation

The SSO token is missing. The token might be missing because of incorrect domain information

Action

Ensure that the application from which you are starting IBM Spectrum Control is in the same TCP/IP domain as the IBM Spectrum Control server. If the application from which you are starting IBM Spectrum Control is in a different TCP/IP domain than the IBM Spectrum Control server, a fully qualified domain name (FQDN) is required for the server. If a FQDN was not provided for the server during installation, you can change the domain name using the web server administrative console.

HWNSS0004E The web server is unavailable and cannot be used for authentication. The Device server can be used for some authentication, but IBM Spectrum Control functions will be limited. To proceed, enter either the common user name that was used to install IBM Spectrum Control or the tpcFileRegistryUser user name and password.

Explanation

The web server provides authentication and single sign-on. If the web server is unavailable, IBM Spectrum Control can only authenticate user names that do not use single sign-on.

Action

Restart the web server and log in to IBM Spectrum Control again.

JSS - Database messages

- [JSS0001I Scheduler service provider started.](#)
- [JSS0002E Scheduler service provider initialization has failed.](#)
- [JSS0003I Scheduler service provider initialization successful.](#)
- [JSS0004I Scheduler service provider shutting down.](#)
- [JSS0005I Scheduler service provider shutdown complete.](#)
- [JSS0006E Unable to connect to repository database in class name.method name.](#)
- [JSS0007E SQL error preparing statement type statement for table table name in class name.method name.](#)
- [JSS0008E SQL error inserting into table table name in class name.method name.](#)
- [JSS0009E SQL error updating table table name in class name.method name.](#)
- [JSS0010E SQL error querying table table name in class name.method name.](#)
- [JSS0011E the job for computer computer name in run run number of job type job creator.job name could not be submitted due to a log table insert error.](#)
- [JSS0012E the job for computer computer name in run run number of job type job creator.job name could not be submitted - the delay limit was exceeded.](#)
- [JSS0013E the job for computer computer name in run run number of job type job creator.job name could not be submitted due to transmit error.](#)
- [JSS0014E run run number of job type job creator.job name could not be started due to SQL update or insert error.](#)
- [JSS0015E The log row for run run number of job type job creator.job name could not be updated due to an SQL error. The run failed/was successful/completed with warnings.](#)
- [JSS0018E The log row for the job for computer computer name in run run number of job type job creator.job name could not be updated due to an SQL error. The job completed with an exit code of exit code.](#)
- [JSS0019E SQL error preparing statement in class name.method name.](#)
- [JSS0020E Unable to find creator and name for schedule schedule ID.](#)
- [JSS0021E Unable to process returned job number job number The job completed with an exit code of exit code.](#)
- [JSS0022E Unknown request type passed to completed job handler.](#)
- [JSS0023E SQL error querying identifier table in class name.method name.](#)
- [JSS0026E Unable to schedule job type job creator.job name because fetch for computer list failed.](#)
- [JSS0027E The obsolete log row for job type job creator.job name could not be deleted due to an SQL error.](#)
- [JSS0046E the job for computer computer name in run run number of job type job creator.job name could not be started due to an agent error.](#)
- [JSS0051E A system error occurred.](#)
- [JSS0052E The agent did not respond.](#)
- [JSS0062W The value for the config file keyword keyword must be a number between number and number. The default \(value\) will be used.](#)
- [JSS0063E Unable to connect to repository database while attempting to delete job log rows from t_run_jobs.](#)
- [JSS0064E SQL error deleting job log rows from t_run_jobs.](#)
- [JSS0066E Unable to issue alert for job type job creator.job name because of following error.](#)
- [JSS0070E Agent could not be reached.](#)
- [JSS0071E Unable to instantiate class class name.](#)
- [JSS0073E the script for computer computer name in run run number of job type job creator.job name could not be submitted due to an error opening script file.](#)
- [JSS0074E The Scheduler service repository connection has failed and it is in auto-restart mode. It will restart automatically when the repository becomes available.](#)
- [JSS0075W The repository connection has been lost. The Scheduler service is terminating and will automatically restart when the repository again becomes available.](#)
- [JSS0076I The Scheduler service is automatically restarting after recovering its repository connection.](#)
- [JSS0094W The job type job named creator.name is scheduled to run once at a time in the past and will not run.](#)
- [JSS2003L The Storage Resource Agent version is not compatible with the server.](#)
- [JSS2006L The agent is disabled.](#)
- [JSS2007W The job for computer computer name in run run number of job type job creator.job name was not started because the Storage Resource Agent is disabled..](#)

JSS0001I Scheduler service provider started.

Explanation

Scheduler service provider started.

JSS0002E Scheduler service provider initialization has failed.

Explanation

Scheduler service provider initialization has failed.

JSS0003I Scheduler service provider initialization successful.

Explanation

Scheduler service provider initialization successful.

JSS0004I Scheduler service provider shutting down.

Explanation

Scheduler service provider shutting down.

JSS0005I Scheduler service provider shutdown complete.

Explanation

Scheduler service provider shutdown complete.

JSS0006E Unable to connect to repository database in class *name.method name*.

Explanation

Unable to connect to repository database in the specified Java program and routine.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

-  [Getting support](#)
-

JSS0007E SQL error preparing *statement type* statement for table *table name* in class *name.method name*.

Explanation

SQL error preparing the specified statement for the specified table in the specified Java program and routine.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

-  [Getting support](#)
-

JSS0008E SQL error inserting into table *table name* in class *name.method name*.

Explanation

SQL error inserting into the specified table in the specified Java program and routine.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0009E SQL error updating table *table name* in *class name.method name*.

Explanation

SQL error updating the specified table in the specified Java program and routine.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0010E SQL error querying table *table name* in *class name.method name*.

Explanation

SQL error querying the specified table in the specified Java program and routine.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0011E the job for computer *computer name* in run *run number* of *job type job creator.job name* could not be submitted due to a log table insert error.

Explanation

The specified job could not be submitted due to a log table insert error.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0012E the job for computer *computer name* in run *run number* of *job type job creator.job name* could not be submitted - the delay limit was exceeded.

Explanation

The specified job could not be submitted - the delay limit was exceeded.

Action

Rerun the job.

JSS0013E the job for computer *computer name* in run *run number* of *job type job creator.job name* could not be submitted due to transmit error.

Explanation

The specified job could not be submitted due to transmit error.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support. Also examine the exit code.

Related reference

-  [Getting support](#)

JSS0014E run *run number* of *job type job creator.job name* could not be started due to SQL update or insert error.

Explanation

The specified job run could not be started due to SQL update or insert error.

Action

Rerun the job.

JSS0015E The log row for run *run number* of *job type job creator.job name* could not be updated due to an SQL error. The run *failed/was successful/completed with warnings*.

Explanation

The log row for the specified job run could not be updated due to an SQL error.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support. Also examine the exit code. If the run failed rerun the job.

Related reference

-  [Getting support](#)

JSS0018E The log row for the job for computer *computer name* in run *run number* of *job type job creator.job name* could not be updated due to an SQL error. The job completed with an exit code of *exit code*.

Explanation

The log row for the specified job run could not be updated due to an SQL error.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support. Also examine the exit code. If greater than four rerun the job.

Related reference

- [Getting support](#)

JSS0019E SQL error preparing statement in *class name.method name*.

Explanation

SQL error preparing statement in the specified Java program and routine.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0020E Unable to find creator and name for schedule *schedule ID*.

Explanation

Restart logic has failed for the above job.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0021E Unable to process returned job number *job number* The job completed with an exit code of *exit code*.

Explanation

This may or may not be an error.

Action

Examine the exit code. If greater than four rerun the job.

JSS0022E Unknown request type passed to completed job handler.

Explanation

An internal error has occurred.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0023E SQL error querying identifier table in *class name.method name*.

Explanation

An SQL error occurred while querying the identifier table in the specified Java program and routine.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0026E Unable to schedule *job type job creator.job name* because fetch for computer list failed.

Explanation

Unable to schedule the specified job because fetch for computer list failed.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0027E The obsolete log row for *job type job creator.job name* could not be deleted due to an SQL error.

Explanation

The obsolete log row for could not be deleted due to an SQL error.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0046E the job for computer *computer name* in run *run number* of *job type job creator.job name* could not be started due to an agent error.

Explanation

The specified job could not be run due to an agent error.

Action

Examine the server log and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0051E A system error occurred.

Explanation

An error occurred contacting the agent

Action

Verify that the agent exists, is running and can be contacted.

JSS0052E The agent did not respond.

Explanation

A response was not received from the agent.

Action

Verify that the agent is running and can be contacted.

JSS0062W The value for the config file keyword *keyword* must be a number between *number* and *number*. The default (*value*) will be used.

Explanation

The value for the specified keyword is out of range. The specified default will be used.

Action

Change the value if the default is not desired.

JSS0063E Unable to connect to repository database while attempting to delete job log rows from t_run_jobs.

Explanation

Unable to connect to repository database while attempting to delete job log rows from t_run_jobs.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0064E SQL error deleting job log rows from t_run_jobs.

Explanation

SQL error deleting job log rows from t_run_jobs.

Action

Examine the associated error text and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0066E Unable to issue alert for job type *job creator.job name* because of following error.

Explanation

Unable to issue alert for the specified because of the following error.

Action

Examine the associated error text and determine what caused the error. If you cannot determine the cause of the error, contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0070E Agent could not be reached.

Explanation

This agent could not be reached.

Action

This message is informational only. Examine other messages to determine the reason this agent can not be reached.

JSS0071E Unable to instantiate class *class name*.

Explanation

Unable to instantiate the specified Java program.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0073E the script for computer *computer name* in run *run number* of job type *job creator.job name* could not be submitted due to an error opening script file.

Explanation

The script for the specified job could not be submitted due to an error opening script file.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

JSS0074E The Scheduler service repository connection has failed and it is in auto-restart mode. It will restart automatically when the repository becomes available.

Explanation

The Scheduler service repository connection has failed and it is in auto-restart mode. It will restart automatically when the repository becomes available.

JSS0075W The repository connection has been lost. The Scheduler service is terminating and will automatically restart when the repository again becomes available.

Explanation

The repository connection has been lost. The Scheduler service is terminating and will automatically restart when the repository again becomes available.

JSS0076I The Scheduler service is automatically restarting after recovering its repository connection.

Explanation

The Scheduler service is automatically restarting after recovering its repository connection.

JSS0094W The *job type* job named *creator.name* is scheduled to run once at a time in the past and will not run.

Explanation

A job is scheduled to run once but the time the job is expected to run has passed. If the message is seen when saving a job there is most likely a time difference between the server and the GUI where the job was saved.

JSS2003L The Storage Resource Agent version is not compatible with the server.

Explanation

The version of the Storage Resource Agent on this machine is not compatible with the version of IBM Spectrum Control running on the server.

Action

Upgrade the Storage Resource Agent on this machine.

JSS2006L The agent is disabled.

Explanation

The Storage Resource Agent on this computer is disabled and will not accept any request from the IBM Spectrum Control server.

Action

Enable the Storage Resource Agent on this computer using the 'Enable' option available for this agent in IBM Spectrum Control Web GUI.

JSS2007W The job for computer *computer name* in run *run number* of job type *job creator.job name* was not started because the Storage Resource Agent is disabled..

Explanation

The specified job was not submitted because the Storage Resource Agent is disabled.

Action

Enable the Storage Resource Agent on this computer using the 'Enable' option available for this agent in IBM Spectrum Control Web GUI and run the job again.

NAD - Storage Resource Agent messages

- [NAD0001I Connecting to hostname using protocol protocol.](#)
- [NAD0002W Connection to hostname failed using protocol protocol: error.](#)
- [NAD0003I Connected to hostname using protocol protocol.](#)
- [NAD0004W Possible cause: provided login information is incorrect.](#)
- [NAD0005E Connection to hostname failed using protocol protocol: error message.](#)
- [NAD0006E Exception thrown for method method name: error message.](#)
- [NAD0007I Closing connection to hostname.](#)
- [NAD0008E Invalid protocol protocol passed to method name.](#)
- [NAD0009E Cannot connect to host name. Remote host is running in a non-global application container.](#)
- [NAD0010E Invalid parameter\(s\) parameter name passed to method name.](#)
- [NAD0011I Validating GUID on remote machine: host name.](#)
- [NAD0012I GUID value validated on remote machine: host name.](#)
- [NAD0013I Installing GUID on remote machine: hostname.](#)
- [NAD0014I GUID successfully installed on remote machine: hostname.](#)
- [NAD0015I GUID not found on remote machine: host name.](#)
- [NAD0016E Could not copy GUID package to remote host: host name.](#)
- [NAD0017E The GUID on remote machine: host_name could not be validated.](#)
- [NAD0018E Command on remote machine: host name failed. Error code = value executing command value.](#)
- [NAD0019E Parameter parameter passed to method is null or 0 length.](#)
- [NAD0020I Host validation succeeded.](#)
- [NAD0021E Host validation failed.](#)
- [NAD0022E Cannot contact remote host due to invalid credentials, check logs for additional information. Host validation failed.](#)
- [NAD0023E There is not enough space on the remote machine. Host validation failed.](#)
- [NAD0024E It was not possible to determine the available space on the remote machine. Host validation failed.](#)
- [NAD0025E The specified directory could not be accessed. Host validation failed.](#)
- [NAD0036E Failed to copy package name to remote machine.](#)
- [NAD0037E Cannot cleanup remote machine directory: directory path.](#)
- [NAD0038E Failed to install agent on host name.](#)
- [NAD0039I Agent successfully installed at install location.](#)
- [NAD0040E Agent install exited with an error code: value.](#)
- [NAD0041E Failed to get agent bundle locations.](#)
- [NAD0042E Exception occurred while retrieving IPs for Data Server machine : host name.](#)
- [NAD0043I Installing agent at install location.](#)
- [NAD0044I Agent registration to Data Server completed successfully.](#)
- [NAD0045I Validation of host name has started.](#)
- [NAD0046E GUID cannot be null value for registration of agent.](#)
- [NAD0047E Exception while reading authentication information.](#)
- [NAD0048E Probe did not start successfully on agent host address, error code returned = value.](#)
- [NAD0049I Running probe on agent host address.](#)
- [NAD0050E Exception occurred while running probe on agent host name : exception message.](#)
- [NAD0051I Successfully started probe on agent host address.](#)
- [NAD0052E Error in receiving remote file name from host address.](#)
- [NAD0053E Exception occurred while receiving remote file name from host address: exception message.](#)
- [NAD0054E The directory install_location is not empty on the remote machine. Host validation failed.](#)
- [NAD0055E Failed to connect to remote host host.](#)
- [NAD0056E Error uninstalling agent at host address:install location.](#)
- [NAD0057E Error uninstalling agent at host address:install location error message.](#)
- [NAD0058I Performing agent upgrade on remote machine host address.](#)
- [NAD0059I Upgrade on host address succeeded.](#)
- [NAD0059E Upgrade on host address failed: error message.](#)
- [NAD0070I Updating Langpacks on remote machine host address.](#)

- [NAD0071W Requested Langpack language pack name not found on local machine.](#)
- [NAD0072I Langpacks updated on remote machine host address.](#)
- [NAD0073E Error updating LangPacks on remote machine host name : error message.](#)
- [NAD0074W No langpacks found on local machine, update not performed.](#)
- [NAD0075E The Agent is already installed on host host_name. Host validation failed.](#)
- [NAD0076E Failed to configure Auto-upgrade feature.](#)
- [NAD0077E Schedule file name file cannot be created because the schedule file is a directory.](#)
- [NAD0078E Cannot create file name.](#)
- [NAD0079E Cannot delete file name.](#)
- [NAD0080W Auto-upgrade feature was already enabled.](#)
- [NAD0081W Auto-upgrade feature was already disabled.](#)
- [NAD0082E Exception has been encountered: exception trace.](#)
- [NAD0083W Following Storage Resource Agent information is invalid agent information.](#)
- [NAD0084W Host location has no match from IBM Spectrum Control cached non-daemon based Storage Resource Agent list. Agent might not be non-Daemon based Storage Resource Agent.](#)
- [NAD0085I Current non-Daemon based Storage Resource Agents in IBM Spectrum Control cache are agents list.](#)
- [NAD0086E Unable to connect to Windows Domain Controller host name.](#)
- [NAD0087E Unable to determine source package home directory.](#)
- [NAD0088E Error deploying package zip file.Error message.](#)
- [NAD0089E Error deleting binary executable file.Error message.](#)
- [NAD0090E No connection exists to remote machine.](#)
- [NAD0091E Error executing binary executable file.Error message.](#)
- [NAD0092E Unable to retrieve remote temporary directory location.](#)
- [NAD0093E Unable to retrieve domains from host name.](#)
- [NAD0094E Unable to retrieve host list for domain name.](#)
- [NAD0095E The platform OS_name is not supported for Storage Resource Agents.](#)
- [NAD0096E Authentication failed due to invalid credentials or insufficient access privileges.](#)
- [NAD0097I Opening connection to hostname.](#)
- [NAD0098I Copying agent files on remote machine...](#)
- [NAD0099I Installing agent on host host name, in install location directory with force option.](#)
- [NAD0100E Agent command did not run successfully on agent host name, error code returned = value.](#)
- [NAD0101I Running agent command on agent host name.](#)
- [NAD0102E Exception occurred while running command on agent host name : error message.](#)
- [NAD0103E Error in deleting file remote data file from host host address.](#)
- [NAD0104E Exception occurred while deleting file remote data file from host host address: exception trace.](#)
- [NAD0105E Failed to lock Agent file.](#)
- [NAD0106E Failed to stop Probe.](#)
- [NAD0107E Failed to stop Agent.](#)
- [NAD0108E Failed to create registry entry on Agent machine.](#)
- [NAD0109E Failed to extract file.](#)
- [NAD0110E Failed to create configuration.](#)
- [NAD0111E Failed to stop Agent service.](#)
- [NAD0112E Agent service already exist.](#)
- [NAD0113E Failed to create Agent service.](#)
- [NAD0114E Failed to start Agent service.](#)
- [NAD0115E Failed to start Probe.](#)
- [NAD0116E Failed to create Agent service.](#)
- [NAD0117E File extraction needs more space.](#)
- [NAD0118E Failed to open archive file.](#)
- [NAD0119E Missing upgrade files.](#)
- [NAD0120E Failed to extract files.](#)
- [NAD0121E Failed to stop probe in upgrade process.](#)
- [NAD0122E Failed to stop Agent in upgrade process.](#)
- [NAD0123E Failed to start Agent in upgrade process.](#)
- [NAD0124E Failed to register Agent to Server.](#)
- [NAD0125E Extraction needs more space.](#)
- [NAD0126E Failed to open archive file.](#)
- [NAD0127I Uninstalling agent from host host address at location path.](#)
- [NAD0128E Failed to start process.](#)
- [NAD0129E Error in opening catalog file.](#)
- [NAD0130E Probe failed.](#)
- [NAD0131E The installation process could not write files to the directory install_location. Host validation failed.](#)
- [NAD0135E The certificate file file_name was not found on host_name.](#)
- [NAD0136E The port number of the agent is invalid.](#)
- [NAD0137E The port number port_number is in use on the remote machine.](#)
- [NAD0138E Invalid character "invalid_character" found in install_location.](#)
- [NAD0139E The User userID does not have sufficient administrator privileges.](#)
- [NAD0140I Cannot ping to host host name.](#)
- [NAD0141E Either the remote host could not be contacted due to invalid credentials or the machine is not reachable. Check logs for more information. Host validation failed.](#)
- [NAD0142E Deployment of Agent failed, error creating startup scripts.](#)
- [NAD0143E Cannot change agent from Daemon based type to non-Daemon based type with force installation, host validation failed.](#)
- [NAD0144E Agent type can not be changed between non-daemon-based and daemon-based with force installation. Host validation failed.](#)
- [NAD0145E Cannot get version information from agent on host.](#)
- [NAD0146E The connection to remote machine failed because the Remote Execution and Access component was unable to create a temporary directory on the remote machine. Remove unneeded ~CSRI* directories in the remote machine's temporary directory.](#)
- [NAD0147E The daemon-based agent on host_address could not be stopped.](#)
- [NAD0148E Daemon based agent on host address failed to start.](#)
- [NAD0149E Runtime files of agent host address are missing or corrupted.](#)
- [NAD0150E Agent is defected.](#)

- [NAD0151E The original install location path cannot be changed.](#)
- [NAD0152E Initialization of TCP/IP failed while creating socket.](#)
- [NAD0153E Agent service could not be stopped at uninstallation time.](#)
- [NAD0154E Invalid server name passed at uninstallation time.](#)
- [NAD0155E Port passed at installation time for agent service is in use.](#)
- [NAD0156E The server host_address cannot be reached because the host name or IP address is not recognized.](#)
- [NAD0157E The server host_name cannot be contacted. The server might be down, unreachable due to network problems, or the SSH credentials might be invalid.](#)
- [NAD0158E GUID value of machine host_name duplicates GUID on agent machine host_address.](#)
- [NAD0160E Agent is already registered with Server on machine host name location path.](#)
- [NAD0161E The install location is not an absolute path.](#)
- [NAD0162E Services script did not run on the host_address server : exception_message.](#)
- [NAD0163I Services files collected from host address machine can be found in local service file archive.](#)
- [NAD0164E Cannot collect service data from the host_address server.](#)
- [NAD0165E Cannot copy the services_file service archive file from the host_address server.](#)
- [NAD0166E Cannot create the local_service directory for storing service archive file.](#)
- [NAD0167I Running scan on agent host address.](#)
- [NAD0168E Exception occurred while running scan on agent host address : exception_message.](#)
- [NAD0169I Successfully started scan on agent host address.](#)
- [NAD0170E Scan did not start successfully on agent host address, error code returned = value.](#)
- [NAD0171E Failed to copy remote file name file to host name machine.](#)
- [NAD0172E Failed to extract remote file name file.](#)
- [NAD0173I Running NetAppImportQuota on agent host address.](#)
- [NAD0174E Exception occurred while running NetAppImportQuota on agent host address : error_message.](#)
- [NAD0175I Successfully started NetAppImportQuota on agent host address.](#)
- [NAD0176E NetAppImportQuota did not start successfully on agent host address, error code returned = value.](#)
- [NAD0180I Installing re-distributable package on.](#)
- [NAD0181I Install of re-distributable package on succeeded.](#)
- [NAD0182E Failed to install re-distributable package on.](#)
- [NAD0183I Validating re-distributable package on host name.](#)
- [NAD0184I Validation of re-distributable package succeeded.](#)
- [NAD0185E Validation of re-distributable package failed.](#)
- [NAD0186I Trying to locate package TIVguid using pkginfo ...](#)
- [NAD0187I Package TIVguid is not installed.](#)
- [NAD0188I Checking TIVguid default install path : path ...](#)
- [NAD0189E Command is not valid.](#)
- [NAD0190E Provided option is not valid.](#)
- [NAD0191E Arguments are not valid.](#)
- [NAD0192E Value for one of the arguments is missing.](#)
- [NAD0193E Localized string missing in message file.](#)
- [NAD0194E Probe process is already running.](#)
- [NAD0195E Failed to open file for writing.](#)
- [NAD0196E Failed to close file.](#)
- [NAD0197E Logfile was not specified.](#)
- [NAD0198E Tracing failed.](#)
- [NAD0199E Cannot start a new Probe process because another one is already running.](#)
- [NAD0200I Validating user on remote machine host name.](#)
- [NAD0201I Validation of user user succeeded.](#)
- [NAD0202E The validation of user credentials for user failed.](#)
- [NAD0203E No tracing.](#)
- [NAD0204E An internal error occurred in the agent process.](#)
- [NAD0205E The provided socket is invalid.](#)
- [NAD0206E Failed to start the service.](#)
- [NAD0207E The file was not found.](#)
- [NAD0208E Not enough memory to run agent.](#)
- [NAD0209E The agent process did not start after the upgrade was finished.](#)
- [NAD0210E The installation directory was not valid when trying to upgrade.](#)
- [NAD0211E The Probe is currently busy.](#)
- [NAD0212E The data file was not found.](#)
- [NAD0213E The exit code is not in the output file.](#)
- [NAD0214E Failed to send the job status.](#)
- [NAD0215E Failed to copy the certificate files.](#)
- [NAD0216E Failed to create directory.](#)
- [NAD0217E Failed to remove directory.](#)
- [NAD0218E Failed to execute the command.](#)
- [NAD0219E Failed to convert wide characters.](#)
- [NAD0220E The installation directory is not valid.](#)
- [NAD0221E No server name was provided.](#)
- [NAD0222E An error occurred while trying to remove entries from the configuration file.](#)
- [NAD0223E Failed to stop Probe process at uninstallation time.](#)
- [NAD0224E Failed to remove registry entry at uninstallation time.](#)
- [NAD0225E Failed to remove service entry at uninstallation time.](#)
- [NAD0226E The agent files are corrupted.](#)
- [NAD0227E Failed to install the GUID.](#)
- [NAD0228E A storage resource agent with a different runtime operation mode, daemon / non-daemon, is already installed at the specified location.](#)
- [NAD0229E Not enough disk space available for installation.](#)
- [NAD0230E An installation is already in progress.](#)
- [NAD0231E Cannot get the server name from the server.](#)
- [NAD0232E The installation directory is not empty.](#)
- [NAD0233E The parameter server name is missing.](#)

- [NAD0234E The parameter server port is missing.](#)
- [NAD0235E The parameter server ip is missing.](#)
- [NAD0236E The parameter agent port is missing.](#)
- [NAD0237E The parameter installation directory is missing.](#)
- [NAD0238E Failed to send the Probe results.](#)
- [NAD0239E Failed to initialize the agent.](#)
- [NAD0240E The service port number is missing.](#)
- [NAD0241E Get data file statistics failed.](#)
- [NAD0242E Get data file read failed.](#)
- [NAD0243E Failed to send the data to the server.](#)
- [NAD0244E Failed to receive the data from the server.](#)
- [NAD0245E The full path was not specified. \(copy file function\)](#)
- [NAD0246E Create file failed. \(copy file function\)](#)
- [NAD0247E Write file failed. \(copy file function\)](#)
- [NAD0248E Open file failed. \(copy file function\)](#)
- [NAD0249E Read file failed. \(copy file function\)](#)
- [NAD0250E The UCS conversion failed.](#)
- [NAD0251E The server connection failed.](#)
- [NAD0252E Please check the OS level.](#)
- [NAD0253W No error message defined for error code: value.](#)
- [NAD0254E Registry entry not found.](#)
- [NAD0255E Insufficient space to copy file file name on remote machine host name in location location.](#)
- [NAD0256E Cannot get the available space on remote machine host name : error message.](#)
- [NAD0257E OS configuration error encountered. Please do a local install with debug set to MAX.](#)
- [NAD0258E Port number is in use on remote machine. Stop process manually or select another port number.](#)
- [NAD0259W Unable to determine Storage Resource Agent version on host . Fabric Discovery will not be invoked.](#)
- [NAD0260I Agent is active.](#)
- [NAD0261I Agent shutdown successfully.](#)
- [NAD0262I Successfully started agent](#)
- [NAD0263W The probe of the agent on host name did not complete in the allocated amount of time.](#)
- [NAD0264I The agent on host name is being probed.](#)
- [NAD0265I After the probe has completed, the new volumes on the host will display in the IBM Spectrum Control GUI.](#)
- [NAD0266I The probe of the agent on host name completed.](#)
- [NAD0267W The probe of the agent on host name completed with return code value.](#)
- [NAD0268I To display newly assigned volumes in the IBM Spectrum Control GUI, start a new probe after the currently running probe completes. Make sure the new probe completes without errors.](#)
- [NAD0269I To display newly assigned volumes in the IBM Spectrum Control GUI, you must start a new probe. Make sure the new probe completes without errors.](#)
- [NAD0270W The Storage Resource Agent on host name is disabled and will not process any requests.](#)
- [NAD0271W The connection to the Storage Resource Agent on host name was not established because the agent is disabled.](#)
- [NAD0272W The connection to the Storage Resource Agent on host name was not established. Retrying using the IP address.](#)
- [NAD0273E The connection to remote machine failed because the Remote Execution and Access component was unable to create a temporary directory on the remote machine. Remove unneeded ~CSRT* directories in the remote machine's temporary directory.](#)
- [NAD0274E An SSH certificate certificate name already exist.](#)
- [NAD0275E Failed to connect to remote host hostname and port. Failed to establish a secure connection.](#)
- [NAD0276E Failed to connect to remote host hostname and port. Failed to establish a secure connection because the SSL handshake failed.](#)
- [NAD0277E Failed to connect to remote host hostname and port. Failed to establish a secure connection because of an invalid SSL key.](#)
- [NAD0278E Failed to connect to remote host hostname and port. Failed to establish a secure connection because the identity of the peer could not be verified.](#)
- [NAD0279E Failed to connect to remote host hostname and port. Failed to establish a secure connection because of an SSL protocol error.](#)
- [NAD0280E The installation failed for the Microsoft VC++ Redistributable package on host. The return code is value.](#)
- [NAD0281E The Storage Resource agent cannot be deployed because of insufficient space or other issues on the target system. The error is: error message.](#)

NAD0001I Connecting to *hostname* using *protocol* protocol.

Explanation

The server is connecting to the Storage Resource Agent at the specified hostname using the specified communication protocol.

NAD0002W Connection to *hostname* failed using *protocol* protocol: *error*.

Explanation

The server failed to connect to the Storage Resource Agent at the specified hostname using the specified communication protocol.

NAD0003I Connected to *hostname* using *protocol* protocol.

Explanation

The server connected to the Storage Resource Agent at the specified hostname using the specified communication protocol.

NAD0004W Possible cause: provided login information is incorrect.

Explanation

The login information might not be entered correctly.

Action

Check if the login information is correct and that the specified user has enough access privileges. You need to check user id and password information entered. If SSH protocol is used and certificate/passphrase is used, then make sure those are valid.

NAD0005E Connection to *hostname* failed using *protocol* protocol: *error message*.

Explanation

The server failed to connect to the Storage Resource Agent at the specified hostname using the specified communication protocol.

NAD0006E Exception thrown for method *method name*: *error message*.

Explanation

An error occurred while processing the specified method.

NAD0007I Closing connection to *hostname*.

Explanation

Closing the connection to the Storage Resource Agent at the specified hostname.

NAD0008E Invalid protocol *protocol* passed to *method name*.

Explanation

The method does not support this protocol.

NAD0009E Cannot connect to *host name*. Remote host is running in a non-global application container.

Explanation

The connection to the specified host cannot be established because the remote host container does not have enough privileges to accept incoming connections.

Action

Please check remote host to which IBM Spectrum Control Server is trying to connect for installation or communication with Storage Resource Agent. You need to start host in global application container mode.

NAD0010E Invalid parameter(s) *parameter name* passed to *method name*.

Explanation

Invalid input to the specified method.

NAD0011I Validating GUID on remote machine: *host name*.

Explanation

The GUID validation process on a host on which Storage Resource Agent is going to be installed is starting. This will check if GUID exists on that host.

NAD0012I GUID *value* validated on remote machine: *host name*.

Explanation

The GUID is pre-installed on remote machine and returns a valid ID. The GUID of the remote machine was successfully validated.

NAD0013I Installing GUID on remote machine: *hostname*.

Explanation

Installing the unique identifier on the specified machine.

NAD0014I GUID successfully installed on remote machine: *hostname*.

Explanation

Unique identifier installed on the specified machine.

NAD0015I GUID not found on remote machine: *host name*.

Explanation

The GUID was not found on the remote machine and will be installed.

NAD0016E Could not copy GUID package to remote host: *host name*.

Explanation

The GUID installation package could not be sent to the remote machine.

Action

Please check if host on which GUID is being installed has sufficient space or enough privileges for user connecting to target host. Other reason could be guid.zip file on IBM Spectrum Control server machine may be corrupt. Please check guid.zip file in <IBM Spectrum Control Server>/data/sra/<OS> directory.

NAD0017E The GUID on remote machine: *host_name* could not be validated.

Explanation

The GUID value could not be validated on the remote host. This could be because the GUID is not installed on the target system or GUID binaries are corrupted.

Action

Check by running "tivguid -show" on the target system, and determine if it returns valid output with the GUID.

NAD0018E Command on remote machine: *host name* failed. Error code = *value* executing command *value*.

Explanation

The installation on the remote machine failed.

NAD0019E Parameter *parameter* passed to *method* is null or 0 length.

Explanation

Invalid input to the specified method.

NAD0020I Host validation succeeded.

Explanation

The host validation process of the host on which Storage Resource Agent is to be installed has completed with success.

NAD0021E Host validation failed.

Explanation

The host validation process of the host on which Storage Resource Agent is to be installed has failed.

Action

Please check Storage Resource Agent deployment logs for more details. You can fix problem specified in error log and re-attempt deployment of Storage Resource Agent (possibly with Force option) to avoid other issues such as non-empty installation directory. Please keep in mind, if force option is used and port is used by other running process then deployment will fail due to port usage.

NAD0022E Cannot contact remote host due to invalid credentials, check logs for additional information. Host validation failed.

Explanation

The remote host validation has failed due to invalid credentials.

Action

Check if the login information is correct and that the specified user has enough access privileges. You need to check user id and password information entered. If SSH protocol is used and certificate/passphrase is used, then make sure those are valid.

NAD0023E There is not enough space on the remote machine. Host validation failed.

Explanation

The remote host validation has failed due insufficient space on the target host of the Storage Resource Agent installation.

Action

Check that the remote host has enough space for the installation of the Storage Resource Agent.

NAD0024E It was not possible to determine the available space on the remote machine. Host validation failed.

Explanation

The remote host validation has failed due to the fact that free space on the host could not be computed.

Action

Check privileges of user used to connect to the target system on which Storage Resource Agent is going to be installed.

NAD0025E The specified directory could not be accessed. Host validation failed.

Explanation

The remote host validation has failed due to the fact that the agent install directory could not be accessed.

Action

Check privileges of the user used to connect to the target system on which the Storage Resource Agent is going to be installed. Check any permissions set on the directory in which the Storage Resource Agent will be installed.

NAD0036E Failed to copy *package name* to remote machine.

Explanation

The specified file could not be copied/extracted to the remote host in specified location.

Action

Please check if target system has enough space, user had sufficient privileges or file to be sent/copied is not corrupt.

NAD0037E Cannot cleanup remote machine directory: *directory path*.

Explanation

For some reasons installation of Storage Resource Agent has failed. At the end of failure it tries to clean up directory on remote host, but it failed to delete it.

Action

First fix the problem reported in failure while attempting to deploy Storage Resource Agent. If you want to re-run the deployment job, try deleting the directory on the target machine . If you cannot delete directory due to some access issues, then attempt to deploy Storage Resource Agent with "Force" option by editing the job definition and selecting the "Force" option.

NAD0038E Failed to install agent on *host name*.

Explanation

The Storage Resource Agent failed to install on specified host.

Action

Please check deployment logs and address issues. Once you address issues, you can re-deploy this job. For any reason, if you cannot delete target directory or any residue exists on target host system, then use "Force" option to re-deploy Storage Resource Agent.

NAD0039I Agent successfully installed at *install location*.

Explanation

The Storage Resource Agent installation has completed successfully.

NAD0040E Agent install exited with an error code: *value*.

Explanation

The Storage Resource Agent installation has failed with the provided error code number.

Action

Most of the time reason for failure is given along with error codes. For additional information, please check installation/user guide for explanation of error codes.

NAD0041E Failed to get agent bundle locations.

Explanation

The Storage Resource Agent source location cannot be found.

Action

Please check <TPCServerLocation>/data/sra/<OS> directory. It should have all deployment packages for Storage Resource Agent. If packages are missing, please copy from CD to that directory.

NAD0042E Exception occurred while retrieving IPs for Data Server machine : *host name*.

Explanation

The internet addresses of the Data Server machine cannot be retrieved.

Action

Contact IBM customer technical support.

NAD0043I Installing agent at *install location*.

Explanation

The Storage Resource Agent will start installing at the specified location on target system.

NAD0044I Agent registration to Data Server completed successfully.

Explanation

The Storage Resource Agent has registered successfully with the Data Server.

NAD0045I Validation of *host name* has started.

Explanation

The validation process for the host is starting.

NAD0046E GUID cannot be null value for registration of agent.

Explanation

The Storage Resource Agent could not be registered since the GUID value is null.

Action

Check the GUID on the target machine. Check by running "tivguid -show" on the target system and verifying that it gives valid output with the GUID. If there are issues with the GUID and no other application is using the GUID on the target system, clean up the GUID and redeploy the Storage Resource Agent on it.

NAD0047E Exception while reading authentication information.

Explanation

Error occurred while reading registration information from host on which Storage Resource Agent is being installed.

Action

Contact IBM customer technical support.

NAD0048E Probe did not start successfully on agent *host address*, error code returned = *value*.

Explanation

The probing process did not start on the host and returned the specified error code.

Action

There may be issue such as missing operating system library. You can extract agent.zip and probe.zip from the CD or <Data Server>/data/sra/<OS> to target system, and run bin/Probe and check log or system out error message.

NAD0049I Running probe on agent *host address*.

Explanation

The probe on the agent is running.

NAD0050E Exception occurred while running probe on agent *host name* : *exception message*.

Explanation

The probe on the host has failed with the provided error message.

Action

Please check probe logs on Storage Resource Agent and Server logs in IBM Spectrum Control Data Server directory.

NAD0051I Successfully started probe on agent *host address*.

Explanation

The probe was started successfully on the Storage Resource Agent.

NAD0052E Error in receiving *remote file name* from *host address*.

Explanation

The remote data file from the host could not be received.

Action

Please check firewall or any other security issues for file access problem.

NAD0053E Exception occurred while receiving *remote file name* from *host address: exception message*.

Explanation

The remote data file from the host could not be received.

Action

Please check firewall or any other security issues for file access problem, please check logs for more information.

NAD0054E The directory *install_location* is not empty on the remote machine. Host validation failed.

Explanation

The target directory for installation of Storage Resource Agent must be empty.

Action

Verify that the install directory for the agent is empty. If you cannot delete/clean-up the installation directory for Storage Resource Agent, use the "force" option, which will bypass checking any residue remaining from an earlier installation.

NAD0055E Failed to connect to remote host *host*.

Explanation

The specified host is down or not reachable via network.

NAD0056E Error uninstalling agent at *host address:install location*.

Explanation

The Storage Resource Agent could not be uninstalled because the connection to the agent could not be made.

Action

If this is daemon based agent, then check if its up and running. If its non-daemon based agent, then please check if authentication information is still valid. You need to check subnet firewall, or firewall on IBM Spectrum Control Server and host on which SRA is being installed. There may be network issue, please check if ping is working.

NAD0057E Error uninstalling agent at *host address:install location error message*.

Explanation

The agent could not be uninstalled due to the reason specified in error message.

Action

Try to address error from specified message. You may need to clean up target host system manually. If you intend to re-install the agent, you can edit the existing deployment job and select the "Force" option.

NAD0058I Performing agent upgrade on remote machine *host address*.

Explanation

The Storage Resource Agent upgrade process has started.

NAD0059I Upgrade on *host address* succeeded.

Explanation

The upgrade of the Storage Resource Agent has succeeded on the host.

NAD0059E Upgrade on *host address* failed: *error message*.

Explanation

The upgrade for the Storage Resource Agent has failed with the specified error message.

Action

Please check error logs on Storage Resource Agent and on Data Server machine. Address errors pointed in logs.

NAD0070I Updating Langpacks on remote machine *host address*.

Explanation

The localized messages on Storage Resource Agents are being upgraded.

NAD0071W Requested Langpack *language pack name* not found on local machine.

Explanation

The package corresponding to localized messages that was requested can not be found on the IBM Spectrum Control Data Server.

Action

Please verify if localized messages (language pack) for IBM Spectrum Control Data Server has been installed. Localized message package for Storage Resource Agent are copied on IBM Spectrum Control Server machine when localized messages for IBM Spectrum Control Data Server gets installed. These localized message packages are used to copy from Data Server to Storage Resource Agent in upgrade of Storage Resource Agent if corresponding localized message is used in upgrade job definition of Storage Resource Agent.

NAD0072I Langpacks updated on remote machine *host address*.

Explanation

The language pack upgrade has succeeded on the remote host.

NAD0073E Error updating LangPacks on remote machine *host name* : *error message*.

Explanation

Error occurred while upgrading localized messages on Storage Resource Agent as indicated in error message.

Action

Please check error logs and address issue.

NAD0074W No langpacks found on local machine, update not performed.

Explanation

The update of the language pack was not performed because there were no previously installed language packs.

Action

Please verify if localized messages (language pack) for IBM Spectrum Control Data Server has been installed. Localized message package for Storage Resource Agent are copied on IBM Spectrum Control Server machine when localized messages for IBM Spectrum Control Data Server gets installed. These localized message packages are used to copy from Data Server to Storage Resource Agent in upgrade of Storage Resource Agent if corresponding localized message is used in upgrade job definition of Storage Resource Agent.

NAD0075E The Agent is already installed on host *host_name*. Host validation failed.

Explanation

IBM Spectrum Control Server has Storage Resource Agent information in its registry.

Action

You cannot install another Agent on the same system. If the Storage Resource Agent on the target host was removed, delete the agent from the IBM Spectrum Control Server. If you want to preserve historical information, you can re-deploy the Agent using the "force" option.

NAD0076E Failed to configure Auto-upgrade feature.

Explanation

The Auto-upgrade feature could not be configured due to an error regarding the schedule file.

Action

Please check existence of the file SCHEDULED_UPGRADES in <TPC>/data directory. This file should be deleted if AutoUpgrade needs to be enabled and vice-versa.

NAD0077E *Schedule file name* file cannot be created because the *schedule file* is a directory.

Explanation

The directory with the same name as the schedule file already exists, so the schedule file cannot be created. This disables AutoUpgrade.

Action

Rename directory which conflicts with schedule file name in <TPC>/data directory and then disable AutoUpgrade.

NAD0078E Cannot create *file name*.

Explanation

The schedule file could not be created as specified in error logs.

Action

Please check logs and take appropriate actions. Please check existence of file SCHEDULED_UPGRADES in <TPC>/data directory. This file should be created if AutoUpgrade needs to be disabled and vice-versa.

NAD0079E Cannot delete *file name*.

Explanation

The schedule file could not be deleted as specified in error logs.

Action

Please check logs and take appropriate actions. Please check existence of file SCHEDULED_UPGRADES in <TPC>/data directory. This file should be deleted if AutoUpgrade needs to be enabled and vice-versa.

NAD0080W Auto-upgrade feature was already enabled.

Explanation

The Auto-upgrade feature cannot be enabled because it is already enabled.

Action

No response needed.

NAD0081W Auto-upgrade feature was already disabled.

Explanation

Cannot disable the Auto-upgrade feature because it is already disabled.

Action

No response needed.

NAD0082E Exception has been encountered: *exception trace*.

Explanation

An exception has been encountered by the program as specified in message.

Action

Please take appropriate action based on error message.

NAD0083W Following Storage Resource Agent information is invalid *agent information*.

Explanation

The specified information about the agent is invalid. This is not expected to cause any problem.

Action

NAD0084W *Host location has no match from IBM Spectrum Control cached non-daemon based Storage Resource Agent list. Agent might not be non-Daemon based Storage Resource Agent.*

Explanation

The non-Daemon Storage Resource Agent could not be found in IBM Spectrum Control Data Server internal cache list while performing upgrade job on those agents.

Action

Restart Data Server and try to run upgrade job again.

NAD0085I *Current non-Daemon based Storage Resource Agents in IBM Spectrum Control cache are agents list .*

Explanation

The agent list of the non-Daemon based SRA currently in the IBM Spectrum Control cache.

NAD0086E *Unable to connect to Windows Domain Controller host name .*

Explanation

Cannot connect to Windows Domain Controller which is used to get list of domains and host list under domain.

Action

Please check network or IBM Spectrum Control Server or Windows Domain controller firewall. Check authentication information used to connect to Windows Domain controller. Make sure "File and Printer Sharing for Microsoft Network" under "Network properties" is enabled. You need to make sure to choose "Network Properties" of a connection which is used to communicate between IBM Spectrum Control Server and Storage Resource Agent.

NAD0087E *Unable to determine source package home directory.*

Explanation

The source package directory required for the domain controller tool cannot be determined.

Action

Please restart IBM Spectrum Control Data Server and check if problem goes away.

NAD0088E *Error deploying package zip file.Error message.*

Explanation

A error has occurred during the deployment of the zip file containing the domain controller tool.

Action

Please check space, access, firewall or any other blocking issues while copying tool to Domain controller.

NAD0089E *Error deleting binary executable file.Error message.*

Explanation

An error has occurred during the removal of the domain controller tool executable.

Action

Please check error logs. You can delete this file manually as well.

NAD0090E No connection exists to remote machine.

Explanation

There are no valid connections to the remote machine.

Action

Please check logs if connection had any issues. You can re-attempt to do what you were doing earlier.

NAD0091E Error executing binary *executable file*.Error message.

Explanation

A error has occurred during the invocation of the domain controller tool executable.

Action

Please check access privileges of user used to connect to domain controller.

NAD0092E Unable to retrieve remote temporary directory location.

Explanation

The temporary directory on domain controller cannot be determined.

Action

Please check access privileges of user used to connect to domain controller.

NAD0093E Unable to retrieve domains from *host name*.

Explanation

The domain list from domain controller cannot be retrieved.

Action

Please check access privileges of user used to connect to domain controller. You may need to contact Domain Controller administrator for more information.

NAD0094E Unable to retrieve host list for *domain name*.

Explanation

The domain list from domain controller cannot be retrieved.

Action

Please check access privileges of user used to connect to domain controller. You may need to contact Domain Controller administrator for more information.

NAD0095E The platform *OS_name* is not supported for Storage Resource Agents.

Explanation

The Storage Resource Agent cannot be installed on the specified platform as it is not one of the supported platforms for Storage Resource Agent.

Action

Check the IBM Spectrum Control support matrix for a list of operating systems on which Storage Resource Agent is supported.

NAD0096E Authentication failed due to invalid credentials or insufficient access privileges.

Explanation

The agent authentication failed due to invalid credentials.

Action

Check user ID and password used for connection between IBM Spectrum Control Server and Storage Resource Agent. If you have used SSH as a connection protocol, and provided a certificate and passphrase, make sure that information is correct.

NAD0097I Opening connection to *hostname*.

Explanation

Opening the connection to the Storage Resource Agent at the specified hostname.

NAD0098I Copying agent files on remote machine...

Explanation

The agent files are currently being copied to the remote host where Storage Resource Agent will get installed.

NAD0099I Installing agent on host *host name*, in *install location* directory with force option.

Explanation

The Storage Resource Agent installation at the specified path will proceed with the force parameter enabled. That means most of the validation will not be performed. If agent is being installed as a daemon (service) and some other running service is using same port as this agent, then installation will fail, as validations are skipped in "Force" option.

NAD0100E Agent command did not run successfully on agent *host name*, error code returned = *value*.

Explanation

The agent on the remote host failed to process the command.

Action

Please check error code in installation/user's guide for more information.

NAD0101I Running agent command on agent *host name*.

Explanation

The agent is now processing the command.

NAD0102E Exception occurred while running command on agent *host name* : *error message*.

Explanation

During the command processing the agent has encountered an error.

Action

Please check error message and take an appropriate action.

NAD0103E Error in deleting file *remote data file* from host *host address*.

Explanation

The data file on the Storage Resource Agent machine cannot be deleted.

Action

Please check access privileges or if any other process is using this file.

NAD0104E Exception occurred while deleting file *remote data file* from host *host address*: *exception trace*.

Explanation

An exception occurred during the deletion of the remote data file on the host.

Action

Please check access privileges or if any other process is using this file.

NAD0105E Failed to lock Agent file.

Explanation

While installing Storage Resource Agent, agent file on remote host cannot be locked.

Action

Please check deployment logs to understand problem. If logs are deleted then re-deploy with "Force" option to leave logs and other files behind.

NAD0106E Failed to stop Probe.

Explanation

The probe cannot be stopped.

Action

Please check deployment logs to understand problem. If logs are deleted then re-deploy with "Force" option to leave logs and other files behind.

NAD0107E Failed to stop Agent.

Explanation

The agent cannot be stopped.

Action

Please check deployment logs to understand problem. If logs are deleted then re-deploy with "Force" option to leave logs and other files behind.

NAD0108E Failed to create registry entry on Agent machine.

Explanation

The registry entry for the agent cannot be created on the agent machine.

Action

It is possible that some registry entries residue exists on target system from an earlier failed installation. Try re-deploying Storage Resource Agent with "Force" option.

NAD0109E Failed to extract file.

Explanation

Installation package files used for deployment of Storage Resource Agent cannot be extracted.

Action

Please check installation package on CD or on IBM Spectrum Control Server at <IBM Spectrum Control Server>/data/sra/<OS> location.

NAD0110E Failed to create configuration.

Explanation

The configuration of Storage Resource Agent has failed in deployment process.

Action

Please check logs and take corrective actions.

NAD0111E Failed to stop Agent service.

Explanation

The agent service did not stop properly.

Action

Please check if there is any issues with service on target machine.

NAD0112E Agent service already exist.

Explanation

The agent service cannot be installed because it already exists.

Action

On target machine service used for Storage Resource Agent already exists. You can cleanup machine manually or try to deploy Storage Resource Agent with "Force" option.

NAD0113E Failed to create Agent service.

Explanation

The agent service could not be created.

Action

Please check agent logs for more information and take corrective action.

NAD0114E Failed to start Agent service.

Explanation

The agent service cannot be started.

Action

Please check agent logs for more information and take corrective action.

NAD0115E Failed to start Probe.

Explanation

The probe could not be started.

Action

Please check agent logs for more information and take corrective action.

NAD0116E Failed to create Agent service.

Explanation

The agent service could not be created.

Action

Please check agent logs for more information and take corrective action.

NAD0117E File extraction needs more space.

Explanation

There is not enough space for the archive to be extracted.

Action

Increase available space on target system in installation directory of Storage Resource Agent and re-attempt installation of Storage Resource Agent.

NAD0118E Failed to open archive file.

Explanation

The archive could not be opened.

Action

Please check for valid archive file by opening it with any zip utility. If file is corrupted then it needs to be replaced from CD or IBM Spectrum Control Server <IBM Spectrum Control Server>/data/sra/<OS> location.

NAD0119E Missing upgrade files.

Explanation

Upgrade of Storage Resource Agent cannot be continued as upgrade files are missing during the upgrade process.

Action

Please check agent logs for more information.

NAD0120E Failed to extract files.

Explanation

Upgrade package files used for deployment of Storage Resource Agent cannot be extracted.

Action

Please check installation package on CD or on IBM Spectrum Control Server at <IBM Spectrum Control Server>/data/sra/<OS> location.

NAD0121E Failed to stop probe in upgrade process.

Explanation

The agent upgrade failed because the running probe could not be stopped.

Action

Please check if there is long running Probe which cannot be stopped for any reason. You can manually stop Probe and then re-attempt the upgrade.

NAD0122E Failed to stop Agent in upgrade process.

Explanation

The agent upgrade failed because the agent could not be stopped.

Action

Please check agent logs for more information.

NAD0123E Failed to start Agent in upgrade process.

Explanation

The agent upgrade failed because the agent could not be started.

Action

Please check agent logs for more information.

NAD0124E Failed to register Agent to Server.

Explanation

The agent could not register with the server.

Action

Please check agent and IBM Spectrum Control Server logs for more information.

NAD0125E Extraction needs more space.

Explanation

The archived files cannot be extracted because there is not enough space.

Action

Please increase space on target directory and re-attempt the failed action.

NAD0126E Failed to open archive file.

Explanation

The archive cannot be opened.

Action

Please check for valid archive file by opening it with any zip utility. If file is corrupted then it needs to be replaced from CD or IBM Spectrum Control Server <IBM Spectrum Control Server>/data/sra/<OS> location.

NAD0127I Uninstalling agent from host *host address* at location *path*.

Explanation

The Storage Resource Agent is being uninstalled.

NAD0128E Failed to start process.

Explanation

The agent process failed to start.

Action

Please check agent logs for more information.

NAD0129E Error in opening catalog file.

Explanation

The catalog file could not be opened.

Action

Please check agent logs for more information. Catalog file is located in <SRA>/nls directory. You can replace it from CD or IBM Spectrum Control Server location. It is in agent.zip file.

NAD0130E Probe failed.

Explanation

The probe failed.

Action

Please check agent logs for more information.

NAD0131E The installation process could not write files to the directory *install_location*. Host validation failed.

Explanation

The host validation failed because files cannot be written to the agent install directory.

Action

Check access privileges for the user account used to deploy Storage Resource Agent.

NAD0135E The certificate file *file_name* was not found on *host_name*.

Explanation

The certificate file could not be found on the host.

Action

Check that the installation package certs.zip has certificate files on CD(if installed locally) or on IBM Spectrum Control Server if deployed remotely.

NAD0136E The port number of the agent is invalid.

Explanation

The valid port range for an agent is between 0 and 65535.

Action

Specify a port number within the valid port range.

NAD0137E The port number *port_number* is in use on the remote machine.

Explanation

The port is already being used by another process on the remote machine.

Action

Either specify a different port which is not being used by any other process or terminate the application that is using it. Use the following commands to determine what process is using the port. In any instance, replace the generic "port_number" with the specific port value.

Windows: netstat -nao | findstr port_number The last column shows the process ID information.

Linux: netstat -nap | grep port_number The last column shows the process ID/Program name information.

AIX: 1. netstat -Aan | grep port_number The hex number in the first column is the address of the protocol control block (PCB) 2. rmsock "addr of PCB" tcpcb This command shows the process that is holding the socket. Note that this command must be run as root.

Solaris: for pid in `ps -ef|grep -v UID|awk '{ print \$2 }'`; do pfiles \$pid 2>/dev/null|awk '/^[0-9]/ {pid_line=\$0} /port: port_number/ {print pid_line; print}'; done The first line shows the process ID information.

HP-UX: for pid in `ps -ef|grep -v UID|awk '{ print \$2 }'`; do pfiles \$pid 2>/dev/null|awk '/^[0-9]/ {pid_line=\$0} /localaddr\port =.*port_number.*remaddr\port =/ {print pid_line; print substr(\$0,0,index(\$0,"remaddr")-1); exit}'; done The first line shows the process ID information.

NAD0138E Invalid character "*invalid_character*" found in *install_location*.

Explanation

The agent install location path contains an invalid character.

Action

Use only supported characters in the directory name. Check the operating system help for supported characters in a directory name.

NAD0139E The User *userID* does not have sufficient administrator privileges.

Explanation

The specified user ID does not have sufficient administrator privileges.

Action

The user needs to have root privileges required for this action on a remote host.

NAD0140I Cannot ping to host *host name*.

Explanation

The host cannot be reached by issuing the ping command.

Action

Please check network or IBM Spectrum Control Server or Storage Resource Agent host machine firewall. Please check if you can run ping successfully from command line.

NAD0141E Either the remote host could not be contacted due to invalid credentials or the machine is not reachable. Check logs for more information. Host validation failed.

Explanation

The host validation failed because the entered credentials are incorrect or the machine cannot be reached.

Action

Use the ping command to check if the machine can be reached. Verify credential information such as user id and password. If you have used SSH as a connection protocol, and provided a certificate and passphrase, make sure that information is correct.

NAD0142E Deployment of Agent failed, error creating startup scripts.

Explanation

The startup scripts for Storage Resource Agent running as daemon on Unix and Linux could not be created.

Action

There could be issue with privileges or scripts with similar name already exist. Please check agent logs for more information.

NAD0143E Cannot change agent from Daemon based type to non-Daemon based type with force installation, host validation failed.

Explanation

The host validation failed because the agent cannot be changed from Daemon to non-Daemon or vice versa.

Action

If you have installed daemon Storage Resource Agent and are trying to convert it to non-daemon agent, it is not possible. Even with the force option, installation will fail. This is true while converting from non-daemon agent to daemon agent as well. If you want to convert from one type to another you need to uninstall and install it again.

NAD0144E Agent type can not be changed between non-daemon-based and daemon-based with force installation. Host validation failed.

Explanation

The host validation failed because the agent cannot be changed between non-daemon and daemon.

Action

It is not possible to change agent type between daemon and non-daemon. Even with the force option, installation will fail. If you want to convert from one type to another you need to uninstall and re-install the Storage Resource Agent.

NAD0145E Cannot get version information from agent on host .

Explanation

Storage Resource Agent version could not be obtained.

Action

Check if there are any connectivity issues with the Storage Resource Agent. If the problem continues, contact IBM support.

Related reference

- [Getting support](#)

NAD0146E The connection to *remote machine* failed because the Remote Execution and Access component was unable to create a temporary directory on the remote machine. Remove unneeded ~CSRI* directories in the remote machine's temporary directory.

Explanation

The connection to the remote machine failed.

Action

Remove unneeded ~CSRI* directories in the remote machine's temporary directory.

NAD0147E The daemon-based agent on *host_address* could not be stopped.

Explanation

The daemon based Storage Resource Agent on the host cannot be stopped.

Action

Check the agent logs and IBM Spectrum Control Server logs for more information.

NAD0148E Daemon based agent on *host address* failed to start.

Explanation

The daemon based Storage Resource Agent on the host cannot be started.

Action

Please check agent logs for more information.

NAD0149E Runtime files of agent *host address* are missing or corrupted.

Explanation

The Storage Resource Agent runtime files on the host are either missing or not behaving as expected.

Action

Restart Storage Resource Agent and check status through GUI. If problem still exists, then re-install Agent with Force option.

NAD0150E Agent is defected.

Explanation

The agent state indicates that it is defected. This may be due to any files missing or corrupted.

Action

Restart Storage Resource Agent and check status through GUI. If the problem still exists, then re-install Agent with Force option.

NAD0151E The original install location *path* cannot be changed.

Explanation

Once the Storage Resource Agent is installed, you cannot change its location.

Action

If you want to change the Storage Resource Agent location, then you must uninstall and re-install the Agent at a new location.

NAD0152E Initialization of TCP/IP failed while creating socket.

Explanation

While creating sockets for the agent the IBM Spectrum Control/IP initialization failed.

Action

Please check available resources on target machine where Storage Resource Agent is being installed. You can see error code in agent logs.

NAD0153E Agent service could not be stopped at uninstallation time.

Explanation

The uninstallation failed because the agent service could not be stopped.

Action

Please check logs for additional information. You can manually stop service and do local uninstallation of Agent. If uninstallation still fails, you can delete agent from IBM Spectrum Control GUI and manually clean up agent files, registry information on target host.

NAD0154E Invalid server name passed at uninstallation time.

Explanation

The uninstallation failed because the server name specified at uninstallation time is not correct.

Action

Please provide proper IBM Spectrum Control server name at uninstallation time. You can see server names in configuration file Agent.config in the <Storage Resource Agent>/config directory.

NAD0155E Port passed at installation time for agent service is in use.

Explanation

The port specified for the Storage Resource Agent is in use by another process.

Action

Please use a port which is not used by any other process on the host.

NAD0156E The server *host_address* cannot be reached because the host name or IP address is not recognized.

Explanation

The host cannot be contacted because the host is unreachable.

Action

Check the host name specified for the target host. Check if the host name can resolve to an IP address from a command line by using commands such as nslookup or ping.

NAD0157E The server *host_name* cannot be contacted. The server might be down, unreachable due to network problems, or the SSH credentials might be invalid.

Explanation

The server cannot be contacted. The server might be down or unreachable due to other connection failures or the server cannot be contacted using the credentials that are provided.

Action

Check the host name that is specified for the server. Check whether the host name can resolve to an IP address from a command line by using commands such as nslookup or ping. Check whether the user and password that are specified are valid. If SSH certificates are used for authentication, check whether the certificate and the passphrase are valid.

NAD0158E GUID *value* of machine *host_name* duplicates GUID on agent machine *host_address*.

Explanation

The GUID value is already in use by another machine.

Action

Regenerate the GUID on one of the machines with the duplicate GUID by running "tivguid -Write -New". You might need to reinstall the Storage Resource Agent after regenerating the GUID.

NAD0160E Agent is already registered with Server on machine *host_name* location *path*.

Explanation

The Storage Resource Agent cannot be registered with the server as another Storage Resource Agent pointing to this IBM Spectrum Control server is already installed at another location on same host.

Action

You can uninstall Storage Resource Agent from other location and then try to install at this new location.

NAD0161E The *install location* is not an absolute path.

Explanation

The specified install location path is not fully qualified.

Action

Provide a fully qualified path for the installation of Storage Resource Agent.

NAD0162E Services script did not run on the *host_address* server : *exception_message*.

Explanation

The script that collects service data about a Storage Resource Agent failed to run on the associated server.

Action

Review the additional error message that was included in the text of this message. Take any actions based on the details of that message. Check the logs of the IBM Spectrum Control server and the Storage Resource agent for more information about the problem.

NAD0163I Services files collected from *host address* machine can be found in *local service file* archive.

Explanation

The service files for the host have been collected and stored in archive file.

NAD0164E Cannot collect service data from the *host_address* server.

Explanation

Service archive files can not be collected from the server on which the Storage Resource agent is installed.

Action

Verify that the server is up and running and that the network is available. Check the logs of the IBM Spectrum Control server and the Storage Resource agent for more information about the problem.

NAD0165E Cannot copy the *services_file* service archive file from the *host_address* server.

Explanation

The service archive file cannot be transferred from the server on which the Storage Resource agent is installed to the IBM Spectrum Control server.

Action

Verify that the Storage Resource agent is up and running and that the network is available. Retry the copy. If the copy fails again, use a command line tool or GUI file manager to locate the directory where the service archive file is located and copy it manually.

NAD0166E Cannot create the *local_service* directory for storing service archive file.

Explanation

The local directory for the service archive file cannot be created.

Action

Check the logs of the Storage Resource agent for more information about why the directory was not created. Use this information to resolve the problem.

NAD0167I Running scan on agent *host address*.

Explanation

The scan of host system has started.

NAD0168E Exception occurred while running scan on agent *host address* : *exception message*.

Explanation

The scan for the host completed with a error message.

Action

Please check agent and IBM Spectrum Control Server logs for more detail.

NAD0169I Successfully started scan on agent *host address*.

Explanation

The scan for the host has successfully started.

NAD0170E Scan did not start successfully on agent *host address*, error code returned = *value*.

Explanation

The scan did not start on the host and returned the provided error code.

Action

Please check the agent logs for more messages and, if necessary, contact IBM Software Support for more information. The default location for the Storage Resource Agent log files is: TPC_installation_directory/agent/log/name_of_server_SRA_communicates_with

NAD0171E Failed to copy *remote file name* file to *host name* machine.

Explanation

Cannot copy input data file from IBM Spectrum Control Server to Storage Resource Agent.

Action

Please check IBM Spectrum Control Data Server and agent logs for more information.

NAD0172E Failed to extract *remote file name* file.

Explanation

Extraction of archive file on Storage Resource Agent machine failed.

Action

Please check installation package on CD or on IBM Spectrum Control Server at <IBM Spectrum Control Server>/data/sra/<OS> location.

NAD0173I Running NetAppImportQuota on agent *host address*.

Explanation

The NetAppImportQuota is running on the Storage Resource Agent.

NAD0174E Exception occurred while running NetAppImportQuota on agent *host address* : *error message*.

Explanation

The NetAppImportQuota command generated an exception as provided in message.

Action

Please check error message on Storage Resource Agent or IBM Spectrum Control Server machine for more information.

NAD0175I Successfully started NetAppImportQuota on agent *host address*.

Explanation

The NetAppImportQuota command has started on the agent.

NAD0176E NetAppImportQuota did not start successfully on agent *host address*, error code returned = *value*.

Explanation

The NetAppImportQuota command could not be started on the agent host generating an exception message.

Action

Please check error message on Storage Resource Agent or IBM Spectrum Control Server machine for more information.

NAD0180I Installing re-distributable package on .

Explanation

Updating remote machine with Visual Studio re-distributable dll package.

NAD0181I Install of re-distributable package on succeeded.

Explanation

Update of Visual studio 2008 Dll's on remote machine succeeded.

NAD0182E Failed to install re-distributable package on .

Explanation

Failed to update remote machine with Visual Studio dll's.

NAD0183I Validating re-distributable package on *host name*.

Explanation

The redistributable package installation is being validated on the Storage Resource Agent host.

NAD0184I Validation of re-distributable package succeeded.

Explanation

The redistributable package installation has been successfully validated on Storage Resource Agent host.

NAD0185E Validation of re-distributable package failed.

Explanation

The redistributable package installation has failed to validate on Storage Resource Agent host.

Action

You can install reinstallation package on Storage Resource Agent machine by using redistrib.zip on CD or from IBM Spectrum Control Server machine from <TPCServer>/data/sra/windows directory. You need to unzip redistrib.zip file and use TivReDist.msi file for installing redistribution package.

NAD0186I Trying to locate package TIVguid using pkginfo ...

Explanation

Trying to locate installation folder of package TIVguid using pkginfo.

NAD0187I Package TIVguid is not installed.

Explanation

Package TIVguid is not installed on the system.

NAD0188I Checking TIVguid default install path : *path* ...

Explanation

Trying to detect existing installation of package TIVguid using the default install path : /opt/tivoli/guid.

NAD0189E Command is not valid.

Explanation

The specified command is not a valid command.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0190E Provided option is not valid.

Explanation

The specified agent option is not valid.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0191E Arguments are not valid.

Explanation

The specified agent command arguments are not valid.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0192E Value for one of the arguments is missing.

Explanation

The value for one of the provided arguments is missing.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0193E Localized string missing in message file.

Explanation

The language specific string is missing from the appropriate message file.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0194E Probe process is already running.

Explanation

Cannot start another probe because an earlier probe process is still running.

Action

You need to wait until earlier probe finishes. If you think this is invalid case, you can restart Storage Resource Agent and try to run probe again. If you still see similar situation, you can restart IBM Spectrum Control Data Server and run probe again.

NAD0195E Failed to open file for writing.

Explanation

The file could not be opened for writing.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0196E Failed to close file.

Explanation

The file could not be closed.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0197E Logfile was not specified.

Explanation

The file in which the log messages should be stored was not specified.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0198E Tracing failed.

Explanation

The trace messages could not be written.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0199E Cannot start a new Probe process because another one is already running.

Explanation

Cannot start a probe because another probe is already running.

Action

You need to wait until earlier probe finishes. If you think this is invalid case, you can restart Storage Resource Agent and try to run probe again. If you still see similar situation, you can restart IBM Spectrum Control Data Server and run probe again.

NAD0200I Validating user on remote machine *host name*.

Explanation

The user credentials for the Storage Resource Agent host are being validated.

NAD0201I Validation of user *user* succeeded.

Explanation

The user credentials of user for the Storage Resource Agent host have been successfully validated.

NAD0202E The validation of user credentials for *user* failed.

Explanation

The user credentials used to access the Storage Resource Agent host are not valid.

Action

Check the user id and password used for connecting to the host. If you used SSH as a connection option, and provided a certificate and passphrase, make sure that the information is correct.

NAD0203E No tracing.

Explanation

There is no tracing possible for the current action.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0204E An internal error occurred in the agent process.

Explanation

The agent process has failed with an internal error.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0205E The provided socket is invalid.

Explanation

The provided socket could not be used.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0206E Failed to start the service.

Explanation

The service could not be started.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information. Once you address problem, please retry starting agent.

NAD0207E The file was not found.

Explanation

The specified file was not found.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0208E Not enough memory to run agent.

Explanation

The host does not have sufficient memory to run the agent.

Action

Please look at the system hardware configuration of Storage Resource Agent, applications running on it and memory usage of different applications on it.

NAD0209E The agent process did not start after the upgrade was finished.

Explanation

The agent process did not start after the upgrade was finished.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information. Please start agent to check if it can be started.

NAD0210E The installation directory was not valid when trying to upgrade.

Explanation

The installation directory was not valid when trying to upgrade.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0211E The Probe is currently busy.

Explanation

The probe process is currently busy and cannot be accessed.

Action

Wait for few minutes for Probe to be responsive. If probe is not responding, try to shut down agent or forcefully stop probe process and restart agent.

NAD0212E The data file was not found.

Explanation

The data file was not found.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0213E The exit code is not in the output file.

Explanation

The exit code is not in the output file.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0214E Failed to send the job status.

Explanation

Failed to send the job status.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information. Communication between IBM Spectrum Control Server and Storage Resource Agent has failed. Please check if IBM Spectrum Control Data Server is up and running.

NAD0215E Failed to copy the certificate files.

Explanation

Failed to copy the certificate files.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information. You need to check network IBM Spectrum Control Server or Storage Resource Agent firewall or any connection issues.

NAD0216E Failed to create directory.

Explanation

Failed to create directory.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information. There may be access privileges issue which needs to be addressed.

NAD0217E Failed to remove directory.

Explanation

Failed to remove directory.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information. There may be access privileges issue which needs to be addressed. You need to check network IBM Spectrum Control Server or Storage Resource Agent firewall or any connection issues.

NAD0218E Failed to execute the command.

Explanation

Failed to execute the command.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0219E Failed to convert wide characters.

Explanation

Failed to convert wide characters.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information. There may be an issue with operating system localization installation for particular language on Storage Resource Agent.

NAD0220E The installation directory is not valid.

Explanation

The installation directory is not valid.

Action

Please make sure valid installation directory has been provided for installation of Storage Resource Agent.

NAD0221E No server name was provided.

Explanation

No server name was provided.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0222E An error occurred while trying to remove entries from the configuration file.

Explanation

An error occurred while trying to remove entries from the configuration file.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0223E Failed to stop Probe process at uninstallation time.

Explanation

Failed to stop Probe process at uninstallation time.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information. Probe process may be busy and could not respond in timely manner. You can forcefully stop probe process and try to uninstall Storage Resource Agent if its not already uninstalled.

NAD0224E Failed to remove registry entry at uninstallation time.

Explanation

Failed to remove registry entry at uninstallation time.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0225E Failed to remove service entry at uninstallation time.

Explanation

Failed to remove service entry at uninstallation time.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information. If Storage Resource Agent has been uninstalled then you need to manually clean up service from host.

NAD0226E The agent files are corrupted.

Explanation

The installed agent files are corrupted.

Action

You can re-install Storage Resource Agent on host with "Force" option. It will just overlay files on top of what was installed.

NAD0227E Failed to install the GUID.

Explanation

The GUID could not be installed.

Action

Please check guid logs for error message and take an appropriate action.

NAD0228E A storage resource agent with a different runtime operation mode, daemon / non-daemon, is already installed at the specified location.

Explanation

The runtime operation mode of a storage resource agent cannot be changed

Action

To deploy a storage resource agent at the same location using a different runtime operation mode, un-install the agent and re-install it with the desired runtime operation mode.

NAD0229E Not enough disk space available for installation.

Explanation

Not enough disk space available for installation.

Action

Please increase available space on Storage Resource Agent host and then re-try installation of agent.

NAD0230E An installation is already in progress.

Explanation

An installation is already in progress.

Action

Please wait until earlier installation process gets completed.

NAD0231E Cannot get the server name from the server.

Explanation

Cannot get the server name from the server.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0232E The installation directory is not empty.

Explanation

The installation directory is not empty.

Action

Please delete installation directory or try to install with "Force" option.

NAD0233E The parameter server name is missing.

Explanation

The parameter server name is missing.

Action

Contact IBM customer technical support.

NAD0234E The parameter server port is missing.

Explanation

The parameter server port is missing.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0235E The parameter server ip is missing.

Explanation

The parameter server ip is missing.

Action

Please provide Server IP address while installing Storage Resource Agent.

NAD0236E The parameter agent port is missing.

Explanation

The parameter agent port is missing.

Action

If you are installing Storage Resource Agent as daemon agent then you need to provide Agent port. Please check usage of Agent.

NAD0237E The parameter installation directory is missing.

Explanation

The parameter installation directory is missing.

Action

Please provide installation directory for installation of Storage Resource Agent.

NAD0238E Failed to send the Probe results.

Explanation

Failed to send the Probe results.

Action

Please check network or IBM Spectrum Control Server or Storage Resource Agent host machine firewall. Please check if you can run ping successfully from command line to IBM Spectrum Control Server and telnet to IBM Spectrum Control Data Server port. Please check IBM Spectrum Control Data Server or Storage Resource Agent logs for additional information.

NAD0239E Failed to initialize the agent.

Explanation

Failed to initialize the agent.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0240E The service port number is missing.

Explanation

The service port number is missing.

Action

If you are installing Storage Resource Agent as daemon agent then you need to provide Agent port. Please check usage of Agent.

NAD0241E Get data file statistics failed.

Explanation

Get data file statistics failed.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0242E Get data file read failed.

Explanation

Get data file read failed.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0243E Failed to send the data to the server.

Explanation

Failed to send the data to the server.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0244E Failed to receive the data from the server.

Explanation

Failed to receive the data from the server.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0245E The full path was not specified. (copy file function)

Explanation

The full path for copying of the file was not specified.

Action

Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0246E Create file failed. (copy file function)

Explanation

During copying of the file the creation of the copy file failed.

Action

There may be access privileges or directory space issue. Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0247E Write file failed. (copy file function)

Explanation

During copying of the file the data could not be written.

Action

There may be access privileges or directory space issue. Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0248E Open file failed. (copy file function)

Explanation

During copying of the file the copied file could not be opened.

Action

There may be access privileges issue. Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0249E Read file failed. (copy file function)

Explanation

During copying of the file the data could not be read.

Action

There may be access privileges issue. Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0250E The UCS conversion failed.

Explanation

The Unicode conversion failed.

Action

Contact IBM customer technical support.

There may be operating system language installation issue. Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0251E The server connection failed.

Explanation

The server connection failed.

Action

Please check network or IBM Spectrum Control Server or Storage Resource Agent host machine firewall. Please check if you can run ping successfully from command line to IBM Spectrum Control Server and telnet to IBM Spectrum Control Data Server port. Please check IBM Spectrum Control Data Server or Storage Resource Agent logs for additional information.

NAD0252E Please check the OS level.

Explanation

The OS level is not correct.

Action

Operating system level is not the one supported for Storage Resource Agent, you cannot use this operating system to install Storage Resource Agent on.

NAD0253W No error message defined for error code: *value*.

Explanation

The specified error code has no associated error message attached.

Action

Please check installation or user's guide for more information.

NAD0254E Registry entry not found.

Explanation

The searched registry entry was not found.

Action

Registry entry used for Storage Resource Agent identification is missing from host. You may need to reinstall agent with force option on host to recover from situation.

NAD0255E Insufficient space to copy file *file name* on remote machine *host name* in location *location*.

Explanation

There is not enough space on the remote host at the specified location path for the file to be copied.

Action

Increase available space on target system and retry failed action.

NAD0256E Cannot get the available space on remote machine *host name* : *error message*.

Explanation

The amount of available space on the remote host could not be obtained.

Action

There may be access privileges issue. Please check Storage Resource Agent and IBM Spectrum Control Server logs for more information.

NAD0257E OS configuration error encountered. Please do a local install with debug set to MAX.

Explanation

The OS is not configured correctly.

Action

Make sure the operating system on this machine has the required software packages installed. Check the software requirements for the operating system and install any missing packages, then try the installation again.

NAD0258E Port *number* is in use on remote machine. Stop process manually or select another port number.

Explanation

The port is already being used by another process on the remote machine.

Action

Process needs to be stopped manually or select another port number.

NAD0259W Unable to determine Storage Resource Agent version on host . Fabric Discovery will not be invoked.

Explanation

Storage Resource Agent version could not be obtained. Since Fabric Discovery is not supported on some early versions of Storage Resource Agent, it will not be invoked.

Action

Check if there are any connectivity issues with the Storage Resource Agent. Also, check the device server message and trace log for more detail. If the problem continues, contact IBM support.

NAD0260I Agent is active.

Explanation

The check to ensure that communication with the agent has completed with success.

NAD0261I Agent shutdown successfully.

Explanation

The agent service was shutdown.

NAD0262I Successfully started agent

Explanation

The agent service was started.

NAD0263W The probe of the agent on *host name* did not complete in the allocated amount of time.

Explanation

The probe did not complete in the allocated amount of time that Server was configured to wait. Probe process continues on agent until completion.

Action

In order to verify that probe completed successfully, please check later the probe log file on the agent host.

NAD0264I The agent on *host name* is being probed.

Explanation

The probe continues on the specified host until completion.

NAD0265I After the probe has completed, the new volumes on the host will display in the IBM Spectrum Control GUI.

Explanation

New volumes might not display immediately in the GUI. This is caused by probes that take longer than expected to complete

Action

Make sure the probe completed successfully. Wait for the new volumes to display.

NAD0266I The probe of the agent on *host name* completed.

Explanation

The probe completed.

NAD0267W The probe of the agent on *host name* completed with return code *value*.

Explanation

The probe of the agent completed with a return code. The Storage Resource agent return codes are documented in the Reference section of the IBM Spectrum Control Information Center: For information about Storage Resource agent return codes, visit <http://pic.dhe.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for Storage Resource agent return codes.

Action

Check the probe log file on the agent host to determine the cause of the return code and resolve the problem. Try the operation again.

NAD0268I To display newly assigned volumes in the IBM Spectrum Control GUI, start a new probe after the currently running probe completes. Make sure the new probe completes without errors.

Explanation

A manual probe is necessary in order to cause the new volumes to display in the IBM Spectrum Control GUI.

Action

If the currently running probe started before the new volumes were assigned to the host, you must start a new probe to display the new volumes in the IBM Spectrum Control GUI. Manually start a new probe and make sure the probe completes successfully.

NAD0269I To display newly assigned volumes in the IBM Spectrum Control GUI, you must start a new probe. Make sure the new probe completes without errors.

Explanation

A manual probe is necessary in order to cause the new volumes to display in the IBM Spectrum Control GUI.

Action

You must start a new probe to display the new volumes in the IBM Spectrum Control GUI. Manually start a new probe and make sure the probe completes successfully.

NAD0270W The Storage Resource Agent on *host name* is disabled and will not process any requests.

Explanation

The Storage Resource Agent specified on this computer is disabled and will not accept any requests from the IBM Spectrum Control server.

Action

Enable the Storage Resource Agent on this computer using the 'Enable' option available for this agent in IBM Spectrum Control Web GUI and submit the request again.

NAD0271W The connection to the Storage Resource Agent on *host name* was not established because the agent is disabled.

Explanation

The Storage Resource Agent specified on this computer is disabled and will not accept any connections from the IBM Spectrum Control server.

Action

Enable the Storage Resource Agent on this computer using the 'Enable' option available for this agent in IBM Spectrum Control Web GUI and submit the request again.

NAD0272W The connection to the Storage Resource Agent on *host name* was not established. Retrying using the IP address.

Explanation

The Storage Resource Agent specified on this computer could not be connected using the fully qualified computer name. Connection is retried using the ip address.

Action

Check if the fully qualified computer name can be resolved through the network using network commands.

NAD0273E The connection to *remote machine* failed because the Remote Execution and Access component was unable to create a temporary directory on the remote machine. Remove unneeded ~CSRI* directories in the remote machine's temporary directory.

Explanation

The connection to the remote machine failed.

Action

Remove unneeded ~CSRI* directories in the remote machine's temporary directory.

NAD0274E An SSH certificate *certificate name* already exist.

Explanation

The SSH certificate could not be copied on IBM Spectrum Control server directory as there is another certificate with this name.

Action

Rename the SSH certificate name and try again.

NAD0275E Failed to connect to remote host *hostname* and *port*. Failed to establish a secure connection.

Explanation

The SSH connection failed due to an unknown SSL error.

Action

Verify the connection is not blocked by a firewall and retry the connection.

NAD0276E Failed to connect to remote host *hostname and port*.
Failed to establish a secure connection because the SSL handshake failed.

Explanation

The SSL connection failed because the IBM Spectrum Control server and the storage resource agent could not negotiate the desired level of security. This could happen if the IBM Spectrum Control server or storage resource agent certificate is not trusted, not valid, or expired.

Action

If you have recently updated the certificate on the IBM Spectrum Control server then make sure you have also replaced the certificate on the storage resource agent. Make sure that the system date on both the IBM Spectrum Control server and storage resource agent machine is within validity date range of both certificates.

NAD0277E Failed to connect to remote host *hostname and port*.
Failed to establish a secure connection because of an invalid SSL key.

Explanation

The SSL connection failed because of a bad SSL key. This is normally caused by a misconfiguration of the SSL certificate and private key.

Action

Verify that the certificates and private keys on the IBM Spectrum Control server and storage resource agent have been configured correctly.

NAD0278E Failed to connect to remote host *hostname and port*.
Failed to establish a secure connection because the identity of the peer could not be verified.

Explanation

The SSL connection failed because the peer was not able to identify itself. This could happen if no certificate is available or the configured cipher suite does not support authorization.

Action

Verify that the certificates on the IBM Spectrum Control server and storage resource agent have been configured correctly. There is no need to investigate the cipher suite because IBM Spectrum Control does not allow this to be configured.

NAD0279E Failed to connect to remote host *hostname and port*.
Failed to establish a secure connection because of an SSL protocol error.

Explanation

The SSL connection failed because of an error in the SSL protocol.

Action

Contact IBM support.

Related reference

- [Getting support](#)

NAD0280E The installation failed for the Microsoft VC++ Redistributable package on *host*. The return code is *value*.

Explanation

The Microsoft VC++ Redistributable package did not install successfully on the remote host.

Action

The return code comes from the Windows operating system. Run the 'net helpmsg' command followed by the return code from a command prompt window to get additional information about the error.

NAD0281E The Storage Resource agent cannot be deployed because of insufficient space or other issues on the target system. The error is: *error message*.

Explanation

This problem might occur when there is insufficient space on the target system, or the target file is read-only or is being used by another application.

Action

Ensure that the target file is not open or in read-only mode, and that enough space is available on the target system. Try the action again.

NAG - Storage Agent Resource messages

- [NAG0008E Please enter a user with administrative privileges for the tree.](#)
- [NAG0009E The user password is required.](#)
- [NAG0030E Please enter a user with administrative privileges on the filer.](#)
- [NAG0108E The Server Name is required.](#)
- [NAG0119E No NDS trees have been located in your installation or they have not yet been assigned a login and password.](#)

NAG0008E Please enter a user with administrative privileges for the tree.

Explanation

A user has not been entered.

Action

Enter a user with administrative privileges for the tree.

NAG0009E The user password is required.

Explanation

The user password is required.

Action

Enter a valid password.

NAG0030E Please enter a user with administrative privileges on the filer.

Explanation

A user has not been entered.

Action

Enter a user with administrative privileges on the filer.

NAG0108E The Server Name is required.

Explanation

A server name is required to manually register a NAS/Netware server.

Action

A server name is required to manually register a NAS/Netware server.

NAG0119E No NDS trees have been located in your installation or they have not yet been assigned a login and password.

Explanation

No NDS trees have been assigned user IDs or passwords.

Action

You must go to Netware Tree Logins screen under Configuration to enter a user ID and password for any discovered NDS trees.

NAS - Storage Agent Resource messages

- [NAS0003E Unable to contact the agent on host *host address*.](#)
- [NAS0004E NetApp quotas: Error gathering list of agents.](#)
- [NAS0005I No filers.](#)
- [NAS0006W No agents available to retrieve quotas from filer.](#)
- [NAS0008I Gathering quotas through agent host address.](#)
- [NAS0009E No agents available.](#)
- [NAS0010E Bad requestData.](#)
- [NAS0011E Not all agents reported back within allotted time.](#)
- [NAS0012W Filer filer : No quotas retrieved.](#)
- [NAS0013E DB error saving/checking quotas.](#)
- [NAS0014E *** INTERNAL ERROR ***.](#)
- [NAS0015I Filer filer : Quota defined for volume *volume* but we have no record of that volume.](#)
- [NAS0016W Quota quota creator.quota name has no consumers.](#)
- [NAS0017E DB error saving/checking quotas \(filer = filer\).](#)
- [NAS0018I NetApp Quota job completed successfully.](#)
- [NAS0019I NetApp Quota job completed with WARNINGS.](#)
- [NAS0020I NetApp Quota job completed with ERRORS.](#)
- [NAS0021I Processed quotas for filer filer.](#)
- [NAS0022E NetApp Quota is still associated with a schedule, unable to delete.](#)
- [NAS0023E NAS server *server name* is already manually registered to a windows domain.](#)
- [NAS0024E Netware server *server name* is already manually registered.](#)
- [NAS0025E An agent is already installed on *server name*. It cannot be manually entered.](#)
- [NAS0028E Network name of filer could not be determined.](#)

NAS0003E Unable to contact the agent on host *host address*.

Explanation

Unable to contact the specified agent.

Action

Examine the associated error messages and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

NAS0004E NetApp quotas: Error gathering list of agents.

Explanation

NetApp quotas: Error gathering list of agents.

Action

Examine the associated error messages and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

NAS0005I No filers.

Explanation

No filers need to be processed.

NAS0006W No agents available to retrieve quotas from *filer*.

Explanation

No agents were found that can be used to process the specified filer.

NAS0008I Gathering quotas through agent *host address*.

Explanation

The quotas will be processed by the specified agent.

NAS0009E No agents available.

Explanation

No agents were found that can be used to process the current filer.

Action

Contact IBM customer technical support.

NAS0010E Bad requestData.

Explanation

An internal error has occurred.

Action

Contact IBM customer technical support.

NAS0011E Not all agents reported back within allotted time.

Explanation

Not all agents reported back within allotted time.

Action

NAS0012W Filer *filer* : No quotas retrieved.

Explanation

No quotas were retrieved from the specified filer.

NAS0013E DB error saving/checking quotas.

Explanation

An SQL error saving/checking quotas.

Action

Examine the associated error messages and determine what caused the error. If unable to determine the cause of the error contact IBM customer technical support.

NAS0014E *** INTERNAL ERROR ***.

Explanation

An internal error has occurred.

Action

Contact IBM customer technical support.

NAS0015I Filer *filer* : Quota defined for volume *volume* but we have no record of that volume.

Explanation

The specified volume will be ignored.

NAS0016W Quota *quota creator.quota name* has no consumers.

Explanation

An internal error has occurred.

Action

Contact IBM customer technical support.

NAS0017E DB error saving/checking quotas (*filer* = *filer*).

Explanation

DB error saving/checking quotas for the specified filer.

Action

Follow the administrator response associated with the associated error messages.

NAS0018I NetApp Quota job completed successfully.

Explanation

NetApp Quota job completed successfully.

NAS0019I NetApp Quota job completed with WARNINGS.

Explanation

NetApp Quota job completed with WARNINGS.

NAS0020I NetApp Quota job completed with ERRORS.

Explanation

The NetApp Quota job completed with ERRORS.

Action

Follow the administrator response associated with the associated error messages.

NAS0021I Processed quotas for filer *filer*.

Explanation

Processed quotas for filer.

NAS0022E NetApp Quota is still associated with a schedule, unable to delete.

Explanation

NetApp Quota is still associated with a schedule, unable to delete.

Action

Follow the administrator response associated with the associated error messages.

NAS0023E NAS server *server name* is already manually registered to a windows domain.

Explanation

The listed server has already been registered from a windows domain.

Action

Each NAS filer can only be registered from one windows computer or domain.

NAS0024E Netware server *server name* is already manually registered.

Explanation

The listed Network server has already been manually registered.

Action

Each Network server can only be manually registered once.

NAS0025E An agent is already installed on *server name*. It cannot be manually entered.

Explanation

An IBM Spectrum Control agent is already running on the entered NAS server.

Action

A NAS server cannot be manually registered if an IBM Spectrum Control agent is already running on it.

NAS0028E Network name of filer could not be determined.

Explanation

The agent assigned to probe the NAS filer could not resolve the network name of the filer.

Action

Check the DNS configuration of the probing agent to ensure it can resolve the network name of the NAS filer.

SAA - Storage Resource Agent - Storage Subsystem messages

- [SAA0001E SYMAPI error error code -- error text.](#)
- [SAA0002E Error connecting to SYMAPI database \(mode = mode\).](#)
- [SAA0003E Error synching Symmetrix ID.](#)
- [SAA0004E Symmetrix Symmetrix ID: Unexpected volume name: volume name.](#)
- [SAA0005E SymDevList\(Symmetrix ID\) failed.](#)
- [SAA0006E SymDiskList\(Symmetrix ID\) failed.](#)
- [SAA0007W Symmetrix Symmetrix ID: SymDiskShow\(Symmetrix volume\) failed.](#)
- [SAA0008I \(Disk disk , Hyper hyper\).](#)
- [SAA0009W No parity hyper found for RAID group 0xRAID group.](#)
- [SAA0010W Symmetrix Symmetrix ID: SymDevShow\(Symmetrix volume\) failed.](#)
- [SAA0011I \(Meta-component number of volume volume\).](#)
- [SAA0012W Symmetrix Symmetrix ID: Volume <Symmetrix volume> not found or already used.](#)
- [SAA0013W Symmetrix Symmetrix ID: Hyper not found, volume remote sequence number \(instance, bus number, target, partition\).](#)
- [SAA0014W Symmetrix Symmetrix ID: SymDevShow\(Symmetrix number\) failed.](#)
- [SAA0015W Symmetrix Symmetrix ID: Volume Symmetrix volume does not contain hyper \(device name, device number, disk ID, hyper number\).](#)
- [SAA0016W Symmetrix Symmetrix ID: Volume Symmetrix name is actually a meta-component of Symmetrix volume.](#)
- [SAA0017W Symmetrix Symmetrix ID: Volume Symmetrix name has no hypers.](#)
- [SAA0018I SYMAPI version: ID.](#)
- [SAA0019E SymShow\(Symmetrix ID\) failed.](#)
- [SAA0020E SymDiscover failed.](#)
- [SAA0021E SymList failed.](#)
- [SAA0022I Storage Subsystem subsystem name \(subsystem alias\) will be probed.](#)
- [SAA0023W Unsupported storage subsystem type: subsystem type.](#)
- [SAA0024E Failed to load SYMAPI.](#)
- [SAA0025W Symmetrix ID: Remote Symmetrix remote ID: Unexpected volume name: name.](#)
- [SAA0026W Symmetrix ID: Volume number \(volume number\) exceeds array capacity \(length\).](#)

SAA0001E SYMAPI error error code -- error text.

Explanation

The specified SYMAPI error occurred.

Action

Follow the administrator response associated with the associated error messages.

SAA0002E Error connecting to SYMAPI database (mode = *mode*) .

Explanation

Error connecting to SYMAPI database.

Action

Follow the administrator response associated with the associated error messages.

SAA0003E Error synching *Symmetrix ID*.

Explanation

There was an error synchronizing the specified Symmetrix volume.

Action

Follow the administrator response associated with the associated error messages.

SAA0004E *Symmetrix Symmetrix ID*: Unexpected volume name: *volume name* .

Explanation

The specified Symmetrix volume has an unexpected volume name. The volume cannot be processed.

Action

Follow the administrator response associated with the associated error messages.

SAA0005E *SymDevList(Symmetrix ID)* failed.

Explanation

SymDevList failed for the specified Symmetrix volume. The volume cannot be processed.

Action

Follow the administrator response associated with the associated error messages.

SAA0006E *SymDiskList(Symmetrix ID)* failed.

Explanation

SymDiskList failed for the specified Symmetrix volume. The volume cannot be processed.

Action

Follow the administrator response associated with the associated error messages.

SAA0007W *Symmetrix Symmetrix ID: SymDiskShow(Symmetrix volume) failed.*

Explanation

SymDevShow failed for the specified Symmetrix volume. The volume cannot be processed.

SAA0008I *(Disk disk , Hyper hyper) .*

Explanation

The specified hyper was not found.

SAA0009W *No parity hyper found for RAID group 0xRAID group.*

Explanation

No parity hyper was found for the specified RAID group.

SAA0010W *Symmetrix Symmetrix ID: SymDevShow(Symmetrix volume) failed.*

Explanation

SymDevShow failed for the specified Symmetrix volume.

SAA0011I *(Meta-component number of volume volume) .*

Explanation

The specified meta-component was not found.

SAA0012W *Symmetrix Symmetrix ID: Volume <Symmetrix volume> not found or already used.*

Explanation

The specified volume was not found or already used for the specified Symmetrix volume.

SAA0013W *Symmetrix Symmetrix ID: Hyper not found, volume remote sequence number (instance, bus number, target, partition) .*

Explanation

The specified hyper was not found for the specified Symmetrix volume.

SAA0014W *Symmetrix Symmetrix ID: SymDevShow(Symmetrix number) failed.*

Explanation

SymDevShow failed for the specified Symmetrix volume.

Action

Follow the administrator response associated with the associated error messages.

SAA0015W *Symmetrix Symmetrix ID: Volume Symmetrix volume does not contain hyper (device name, device number, disk ID, hyper number).*

Explanation

The specified Symmetrix volume does not contain the specified hyper.

SAA0016W *Symmetrix Symmetrix ID: Volume Symmetrix name is actually a meta-component of Symmetrix volume.*

Explanation

The specified Symmetrix volume has already been processed.

SAA0017W *Symmetrix Symmetrix ID: Volume Symmetrix name has no hypers.*

Explanation

The specified Symmetrix volume has no hypers.

SAA0018I *SYMAPI version: ID.*

Explanation

The specified SYMAPI version is being used.

SAA0019E *SymShow (Symmetrix ID) failed.*

Explanation

Processing for the specified device has failed.

Action

Follow the administrator response associated with the associated error messages.

SAA0020E *SymDiscover failed.*

Explanation

An internal error has occurred.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

SAA0021E SymList failed.

Explanation

An internal error has occurred.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

SAA0022I Storage Subsystem *subsystem name (subsystem alias)* will be probed.

Explanation

The specified storage subsystem will be probed.

SAA0023W Unsupported storage subsystem type: *subsystem type*.

Explanation

The specified storage subsystem type cannot be processed.

Action

Follow the administrator response associated with the associated error messages.

SAA0024E Failed to load SYMAPI.

Explanation

An internal error has occurred.

Action

Contact IBM customer technical support.

Related reference

- [Getting support](#)

SAA0025W Symmetrix ID: Remote Symmetrix *remote ID*: Unexpected volume name: *name*.

Explanation

The specified Symmetrix volume cannot be processed.

SAA0026W Symmetrix ID: Volume number (*volume number*) exceeds array capacity (*length*).

Explanation

The specified Symmetrix volume exceeds the product array capacity and cannot be processed.

VPLG - VASA provider messages

- [VPLUG0001E The connection was refused because the IBM Spectrum Control server is not available at the specified host and port.](#)
- [VPLUG0002E Unable to find the IBM Spectrum Control server configuration information.](#)
- [VPLUG0003E Unable to access the IBM Spectrum Control server configuration information.](#)
- [VPLUG0004E The service class service_class_name does not exist.](#)
- [VPLUG0005E An error was received from the IBM Spectrum Control server: Key:message_key; Message:message_text; Status code: status_code](#)
- [VPLUG0006E An error was received from the IBM Spectrum Control server: message](#)
- [VPLUG0007E An error was encountered while retrieving the list of capacity pools.](#)
- [VPLUG0008E Authentication to the IBM Spectrum Control server failed.](#)
- [VPLUG0009E The specified host is not known.](#)
- [VPLUG0010E The specified host cannot establish a secure connection.](#)
- [VPLUG0011E An error was encountered when retrieving the list of service classes.](#)
- [VPLUG0012E An error was encountered while provisioning block storage.](#)
- [VPLUG0013E An error was encountered while provisioning file storage.](#)
- [VPLUG0014E The provisioning task timed out.](#)
- [VPLUG0015E An error was encountered while retrieving the status of the provisioning task.](#)
- [VPLUG0016E An error was encountered while retrieving the list of WWPNs.](#)
- [VPLUG0017E An error was encountered when retrieving the file storage NFS options.](#)
- [VPLUG0018E An error was encountered while checking administrator privilege.](#)
- [VPLUG0019E An error was encountered while retrieving fabric port information.](#)
- [VPLUG0020E An error was encountered while retrieving the IBM Spectrum Control port.](#)
- [VPLUG0021E An error occurred while retrieving virtual disk information from IBM Spectrum Control.](#)
- [VPLUG0022E An error was encountered while retrieving volume performance information from IBM Spectrum Control.](#)
- [VPLUG0023E An error was encountered while retrieving storage devices list from vCenter.](#)
- [VPLUG0024E An error occurred while retrieving virtual machine virtual disk information from vCenter.](#)
- [VPLUG0025E An error was encountered while retrieving available datastore names from vCenter.](#)
- [VPLUG0026E An error was encountered while retrieving the Lun WWN.](#)
- [VPLUG0027E An error was received from the IBM Spectrum Control server.](#)
- [VPLUG0028E An error was received from the IBM Spectrum Control server: Key:message_key; Message:message_text](#)
- [VPLUG0029E An error was encountered while retrieving roles for current user.](#)
- [VPLUG0030E The user {0} does not have the minimum role \(Administrator, Monitor or External application\) needed to access IBM Spectrum Control.](#)
- [VPLUG0031E Registration of the IBM Spectrum Control VASA provider could not be completed.](#)
- [VPLUG0032E The IBM Spectrum Control VASA provider is already registered for server server_name. Register IBM Spectrum Control VASA provider manually to update credentials.](#)
- [VPLUG0033E One or more third party VASA provider\(s\) are already registered with the vCenter. IBM Spectrum Control VASA provider was not registered. Register IBM Spectrum Control VASA provider manually.](#)
- [VPLUG0034E Automatic registration of IBM Spectrum Control VASA provider is not supported for vCenter server version 5.0 and previous.](#)
- [VPLUG0035E Error encountered while saving IBM Spectrum Control server configuration. Incorrect value for user name or password.](#)
- [VPLUG0036E Error encountered while saving IBM Spectrum Control server configuration. Invalid host name message.](#)
- [VPLUG0037E Error encountered while saving IBM Spectrum Control server configuration. Invalid port message.](#)
- [VPLUG0038E Error encountered while loading IBM Spectrum Control server configuration. Incorrect value for user name or password.](#)
- [VPLUG0039E Error encountered while loading IBM Spectrum Control server configuration. Invalid port message.](#)
- [VPLUG0040E An error was encountered while the server configuration information for IBM Spectrum Control was saved. The current session is invalid.](#)
- [VPLUG0041E An error was encountered while checking administrator privilege. The current session is invalid.](#)

VPLUG0001E The connection was refused because the IBM Spectrum Control server is not available at the specified host and port.

Explanation

The IBM Spectrum Control server is not responding. If the specified host and port are correct, the server might not be running properly or there might be problems with the network connections.

Action

Verify that the IBM Spectrum Control server is running at the specified host and port and that the network is not experiencing connection problems.

VPLUG0002E Unable to find the IBM Spectrum Control server configuration information.

Explanation

The configuration information has not been saved.

Action

Use the Administrator interface of the IBM Spectrum Control action to create the IBM Spectrum Control server configuration information.

VPLUG0003E Unable to access the IBM Spectrum Control server configuration information.

Explanation

The configuration file may be corrupted.

Action

Contact the system administrator for further assistance.

VPLUG0004E The service class *service_class_name* does not exist.

Explanation

IBM Spectrum Control cannot find the specified service class. The service class might have been deleted.

Action

Verify that the service class exists in IBM Spectrum Control or select a different service class.

VPLUG0005E An error was received from the IBM Spectrum Control server: *Key:message_key; Message:message_text; Status code: status_code*

Explanation

An error was received from the server.

Action

No action is required.

VPLUG0006E An error was received from the IBM Spectrum Control server: *message*

Explanation

An error was received from the server.

Action

No action is required.

VPLUG0007E An error was encountered while retrieving the list of capacity pools.

Explanation

An unknown error was encountered while the server was retrieving the list of capacity pools.

Action

No action is required.

VPLUG0008E Authentication to the IBM Spectrum Control server failed.

Explanation

Authentication to the server failed. This might be caused by an incorrect user name or password.

Action

Verify that the user name and password are entered correctly in the Configuration panel.

VPLUG0009E The specified host is not known.

Explanation

The specified host cannot be contacted. This might be caused by an incorrect host name in the Configuration panel, or by network connection problems.

Action

Verify that the host name is correct in the Configuration panel. Verify that the network is not experiencing problems.

VPLUG0010E The specified host cannot establish a secure connection.

Explanation

The server encountered an error while trying to establish a secure connection. This might be caused by an incorrect port specification in the Configuration panel, or by network connection problems.

Action

Verify that the port is correct in Configuration panel. Verify that the network is not experiencing problems.

VPLUG0011E An error was encountered when retrieving the list of service classes.

Explanation

An unknown error was encountered when retrieving the list of service classes.

Action

No action required.

VPLUG0012E An error was encountered while provisioning block storage.

Explanation

The provisioning block storage task failed due to unknown error.

Action

No action required.

VPLUG0013E An error was encountered while provisioning file storage.

Explanation

The provisioning file storage task failed due to unknown error.

Action

No action required.

VPLUG0014E The provisioning task timed out.

Explanation

The provisioning task did not complete in the specified timeout period.

Action

Increase the time out period in the Configuration panel and try the operation again.

VPLUG0015E An error was encountered while retrieving the status of the provisioning task.

Explanation

An unknown error occurred while retrieving the status of the provisioning task.

Action

No action is required.

VPLUG0016E An error was encountered while retrieving the list of WWPNS .

Explanation

An unknown error occurred while retrieving the list of WWPNS.

Action

No action is required.

VPLUG0017E An error was encountered when retrieving the file storage NFS options.

Explanation

An unknown error was encountered when retrieving the file storage NFS options.

Action

No action required.

VPLUG0018E An error was encountered while checking administrator privilege.

Explanation

An unknown error was encountered while checking administrator privilege.

Action

No action required.

VPLUG0019E An error was encountered while retrieving fabric port information.

Explanation

An unknown error was encountered while retrieving fabric port information.

Action

No action required.

VPLUG0020E An error was encountered while retrieving the IBM Spectrum Control port.

Explanation

An unknown error was encountered while retrieving the IBM Spectrum Control port.

Action

No action required.

VPLUG0021E An error occurred while retrieving virtual disk information from IBM Spectrum Control.

Explanation

An unknown error was encountered while retrieving virtual disk information from IBM Spectrum Control.

Action

Check IBM Spectrum Control logs for related error messages.

VPLUG0022E An error was encountered while retrieving volume performance information from IBM Spectrum Control.

Explanation

Retrieving volume performance information from IBM Spectrum Control could not be completed.

Action

Check IBM Spectrum Control logs for related error messages.

VPLUG0023E An error was encountered while retrieving storage devices list from vCenter.

Explanation

Retrieving volume performance information from vCenter could not be completed.

Action

Check vCenter logs for related error messages.

VPLUG0024E An error occurred while retrieving virtual machine virtual disk information from vCenter.

Explanation

Retrieval of virtual disk information from vCenter could not be completed.

Action

Check vCenter logs for related error messages.

VPLUG0025E An error was encountered while retrieving available datastore names from vCenter.

Explanation

Retrieving datastore names from vCenter could not be completed.

Action

Check vCenter logs for related error messages.

VPLUG0026E An error was encountered while retrieving the Lun WWN.

Explanation

An unknown error occurred while retrieving the Lun WWN.

Action

Check vCenter logs for related error messages.

VPLUG0027E An error was received from the IBM Spectrum Control server.

Explanation

An unknown error was received from the server.

Action

No action is required.

VPLUG0028E An error was received from the IBM Spectrum Control server: *Key:message_key; Message:message_text*

Explanation

An error was received from the server.

Action

No action is required.

VPLUG0029E An error was encountered while retrieving roles for current user.

Explanation

An unknown error was encountered while retrieving roles for current user.

Action

Check IBM Spectrum Control logs for related error messages.

VPLUG0030E The user {0} does not have the minimum role (Administrator, Monitor or External application) needed to access IBM Spectrum Control.

Explanation

A user should have one of the following roles: Administrator, Monitor or External application to be able to access the Plugin and VASA functionality.

Action

Login as a user with atleast the minimal role in IBM Spectrum Control.

VPLUG0031E Registration of the IBM Spectrum Control VASA provider could not be completed.

Explanation

An unexpected error occurred while attempting to register the IBM Spectrum Control VASA provider.

Action

Check IBM Spectrum Control logs for related error messages.

VPLUG0032E The IBM Spectrum Control VASA provider is already registered for server *server_name*. Register IBM Spectrum Control VASA provider manually to update credentials.

Explanation

Automatic update of IBM Spectrum Control VASA provider is not supported if the provider is already registered for this server.

Action

Remove the IBM Spectrum Control VASA provider and register it again manually to update the user name or password. For more information visit: <http://pic.dhe.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for "Registering VASA provider".

VPLUG0033E One or more third party VASA provider(s) are already registered with the vCenter. IBM Spectrum Control VASA provider was not registered. Register IBM Spectrum Control VASA provider manually.

Explanation

If one or more third party VASA provider(s) are already registered with the vCenter, you must register IBM Spectrum Control VASA provider manually.

Action

Register IBM Spectrum Control VASA provider manually. For more information visit: <http://pic.dhe.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for "Registering VASA provider".

VPLUG0034E Automatic registration of IBM Spectrum Control VASA provider is not supported for vCenter server version 5.0 and previous.

Explanation

vCenter server version 5.1 or later is required to automatically register IBM Spectrum Control VASA provider.

Action

Register the IBM Spectrum Control VASA provider manually. For more information visit: <http://pic.dhe.ibm.com/infocenter/tivihelp/v59r1/index.jsp> and search for "Registering VASA provider".

VPLUG0035E Error encountered while saving IBM Spectrum Control server configuration. Incorrect value for user name or password.

Explanation

Error encountered while saving IBM Spectrum Control server configuration.

Action

Save the IBM Spectrum Control server configuration using correct user name and/or password.

VPLUG0036E Error encountered while saving IBM Spectrum Control server configuration. Invalid host name *message*.

Explanation

Error encountered while saving IBM Spectrum Control server configuration.

Action

Save the IBM Spectrum Control server configuration using valid host name.

VPLUG0037E Error encountered while saving IBM Spectrum Control server configuration. Invalid port *message*.

Explanation

Error encountered while saving IBM Spectrum Control server configuration.

Action

Save the IBM Spectrum Control server configuration using valid port. Default value is 9569.

VPLUG0038E Error encountered while loading IBM Spectrum Control server configuration. Incorrect value for user name or password.

Explanation

Error encountered while loading IBM Spectrum Control server configuration.

Action

Save the IBM Spectrum Control server configuration using correct user name and/or password.

VPLUG0039E Error encountered while loading IBM Spectrum Control server configuration. Invalid port *message*.

Explanation

Error encountered while loading IBM Spectrum Control server configuration.

Action

Save the IBM Spectrum Control server configuration using valid port. Default value is 9569.

VPLUG0040E An error was encountered while the server configuration information for IBM Spectrum Control was saved. The current session is invalid.

Explanation

The current session is invalid.

Action

Log out of the vSphere Web Client and then log back in.

VPLUG0041E An error was encountered while checking administrator privilege. The current session is invalid.

Explanation

The current session is invalid.

Action

Log out of the vSphere Web Client and then log back in.

Publications

A number of publications are provided with IBM Spectrum® Control.

The following section describes how to access these publications online.

- **Accessing publications online**
Information about installing, configuring, using, upgrading, and uninstalling IBM Spectrum Control and related products is available online.
- **IBM Redbooks**
The IBM® Redbooks® are publications about specialized topics.
- **Languages**
Non-English publications are available from IBM Documentation, which is available in certain non-English languages. It is displayed in the language that is appropriate for the browser locale setting.
- **Providing feedback about publications**
Your feedback is important to help IBM provide the highest quality information.

Accessing publications online

Information about installing, configuring, using, upgrading, and uninstalling IBM Spectrum® Control and related products is available online.

Use [Table 1](#) to view and download these publications. Translated documents are available for some products.

Table 1. Locations of publications for IBM Spectrum Control and related products

Product	Online location
IBM® Cognos® Analytics	https://www.ibm.com/docs/en/cognos-analytics
IBM® Db2® Database	https://www.ibm.com/docs/en/db2
IBM FlashSystem® 5000, IBM FlashSystem 5100, IBM FlashSystem 5200 and IBM Storwize® V5000E	https://www.ibm.com/docs/en/v5200
IBM FlashSystem 7300, IBM FlashSystem 7200 and IBM Storwize V7000	https://www.ibm.com/docs/en/flashsystem-7x00
IBM FlashSystem 900	https://www.ibm.com/docs/en/flashsystem-900/1.6.1
IBM FlashSystem 9500, IBM FlashSystem 9200 and IBM FlashSystem 9100	https://www.ibm.com/docs/en/flashsystem-9x00
IBM FlashSystem V9000	https://www.ibm.com/docs/en/flashsystem-v9000
IBM FlashSystem A9000	https://www.ibm.com/docs/en/flashsystem-a9000
IBM FlashSystem A9000R	https://www.ibm.com/docs/en/flashsystem-a9000r
IBM Spectrum Accelerate	https://www.ibm.com/docs/en/spectrum-accelerate
IBM Storage Insights	https://www.ibm.com/docs/en/storage-insights
IBM Spectrum Virtualize as Software Only	https://www.ibm.com/docs/en/spectrumvirtualsoftw
IBM Spectrum Virtualize for Public Cloud	https://www.ibm.com/docs/en/spectrumvirtualizecl
IBM Spectrum Scale (General Parallel File System)	https://www.ibm.com/docs/en/gpfs
IBM Storwize V3500	https://www.ibm.com/docs/en/v3500/7.8.1
IBM Storwize V3700	https://www.ibm.com/docs/en/v3700
IBM FlashSystem 5000, IBM FlashSystem 5100, IBM FlashSystem 5200, and IBM Storwize V5000E	https://www.ibm.com/docs/en/v5200
IBM Storwize V7000 Unified	https://www.ibm.com/docs/en/flashsystem-v7000u
IBM System Storage® DS8000®	https://www.ibm.com/docs/en Go to the preceding website and search for the DS8000 series system that you want.
IBM System Storage N series	https://www.ibm.com/docs/en/n-series?topic=nseries/kc_welcome_nseries.html
IBM SAN Volume Controller	https://www.ibm.com/docs/en/sanvolumecontroller/8.3.1
IBM XIV® Storage System	https://www.ibm.com/docs/en/xiv-storage-system
IBM WebSphere® Application Server	https://www.ibm.com/docs/en/was
IBM WebSphere Application Server Liberty	https://www.ibm.com/docs/en/was-liberty/base

IBM Redbooks

The IBM® Redbooks® are publications about specialized topics.

You can order publications through your IBM representative or the IBM branch office serving your locality. You can also search for and order books of interest to you by visiting the IBM Redbooks home page at <http://www.redbooks.ibm.com>.

Languages

Non-English publications are available from IBM® Documentation, which is available in certain non-English languages. It is displayed in the language that is appropriate for the browser locale setting.

When a locale does not have a non-English version, the information is displayed in English, which is the default language. Non-English PDFs are available when the information is translated.

Providing feedback about publications

Your feedback is important to help IBM® provide the highest quality information.

To provide comments or suggestions about the product documentation, go to the IBM Knowledge Center for IBM Spectrum® Control: <https://www.ibm.com/docs/en/spectrum-control>. Click Feedback at the bottom of any page and complete the form.

Legal notices

This information was developed for products and services offered in the U.S.A. This material may be available from IBM® in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

About this task

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
224A/101
11400 Burnet Road
Austin, TX 78758
U.S.A*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE: This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

- [Trademarks](#)
- [Privacy policy considerations](#)

Trademarks

IBM®, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web on the Copyright and trademark information page at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft is a trademark of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Privacy policy considerations

For information about privacy policy considerations, see IBM®'s Privacy Policy at <https://www.ibm.com/privacy>.

Glossary

This glossary provides terms and definitions for IBM Spectrum® Control.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

[A](#) [B](#) [C](#) [D](#) [E](#) [F](#) [G](#) [H](#) [I](#) [J](#) [K](#) [L](#) [M](#) [N](#) [O](#) [P](#) [R](#) [S](#) [T](#) [U](#) [V](#) [W](#) [Z](#)

A

activation key

See [license key](#).

active management server

A management server from which the storage environment can be monitored and managed. The active management server replicates its database to the standby server.

advisory lock

A type of lock that a process holds on a region of a file that signals any other process to not use or lock the region or an overlapping region. Other processes are not forced to comply.

allocatable extent limit

A maximum total capacity for the system. The allocatable extent limit is calculated from pool extent sizes.

application key

See [license key](#).

array

An ordered collection, or group, of physical devices (disk drive modules) that are used to define logical volumes or devices. An array is a group of drives designated to be managed with a Redundant Array of Independent Disks (RAID).

asynchronous replication

A type of replication in which control is given back to the application as soon as the write operation is made to the source volume. Some time later, the write operation is made to the target volume. See also [synchronous replication](#).

audit log

An unalterable record of all commands or user interactions that are issued to the system.

authenticated user

A user who has logged in to the system with a valid account (user ID and password).

authentication

The mechanism by which a system determines what permissions a particular authenticated user has to access specific resources or actions. See also [authorization](#).

authorization

The mechanism by which a system determines what permissions a particular authenticated user has to access specific resources or actions. See also [authentication](#).

authorization code

An alphanumeric code generated for administrative functions, such as password resets or two-factor authentication bypass.

available capacity

The amount of usable capacity that is not yet used in a system, pool, array, or MDisk.

B

block storage

A unit of data storage on a device.

business continuity

The capability of a business to withstand outages and to operate mission-critical services normally and without interruption in accordance with predefined service-level agreements.

C

cache

Storage or memory that is used to improve access times to instructions, data, or both. For example, data that resides in cache memory is normally a copy of data that resides elsewhere in slower, less expensive storage, such as on a disk or on another network node.

cache eviction

A process by which data associated with a file is removed from the cache system. The data is removed either by using a Least Recently Used (LRU) algorithm when configured General Parallel File System (GPFS) hard or soft quota limits are exceeded or by issuing a command. When referenced again in the cache system, the data that is associated with the file is retrieved from the home system.

caching I/O group

The I/O group in the system that performs the cache function for a volume.

call home

A communication link established between a product and a service provider. The product can use this link to place a call to a service provider when it requires service. With access to the machine, service personnel can perform service tasks, such as viewing error and problem logs or initiating trace and dump retrievals.

capacity

The amount of data that can be contained on a storage medium.

capacity recycling

The amount of provisioned capacity that can be recovered without causing stress or performance degradation. This capacity identifies the amount of resources that can be reclaimed and provisioned to other objects in an environment.

capacity threshold

The percent of total usable physical capacity that used capacity must exceed before a notification is sent. See also [total usable physical capacity](#).

certificate

A digital document that binds a public key to the identity of the certificate owner, thereby enabling the certificate owner to be authenticated. A certificate is issued by a certificate authority and is digitally signed by that authority.

change volume

A volume that is used in Global Mirror that holds earlier consistent revisions of data when changes are made.

child pool

A user-defined capacity that is formed from capacity that is defined either in another pool or a system. See also [parent pool](#).

CIFS

See [Common Internet File System](#).

CIM

See [Common Information Model](#).

CIM agent

The code that consists of common building blocks that can be used instead of proprietary software or device-specific programming interfaces to manage devices that are compliant with the Common Information Model (CIM).

CIM object manager (CIMOM)

The common conceptual framework for data management that receives, validates, and authenticates the CIM requests from the client application. It then directs the requests to the appropriate component or service provider.

CIMOM

See [CIM object manager](#).

CKD

See [count key data](#).

CKD record

See [count-key-data record](#).

CLI

See [command-line interface](#).

client

A software program or computer that requests services from a server. See also [host, server](#).

cloud account

An agreement with a cloud service provider to use storage or other services at that service provider. Access to the cloud account is granted by presenting valid credentials.

cluster

1. A collection of complete systems that work together to provide a single, unified computing capability.
2. A group of computers and other resources that operate together as a single system.
3. A loosely coupled collection of independent systems (or nodes) organized into a network for the purpose of sharing resources and communicating with each other.
4. In IBM® System Storage DS8000®, a partition capable of performing all DS8000 series functions. With two clusters in the DS8000 storage unit, any operational cluster can take over the processing of a failing cluster.
5. In Storwize® V7000, a pair of nodes that provides a single configuration and service interface.

command-line interface (CLI)

A computer interface in which the input and output are text based.

Common Information Model (CIM)

An implementation-neutral, object-oriented schema for describing network management or systems management information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

Common Internet File System (CIFS)

A protocol that manages shared, remote file access for applications to files, printers, serial ports, and so on over a TCP/IP network.

community name

The part of an SNMP message that represents a password-like name and that is used to authenticate the SNMP message.

compression

A function that removes repetitive characters, spaces, strings of characters, or binary data from the data being processed and replaces characters with control characters. Compression reduces the amount of storage space that is required for data.

compute node

An independent machine that contains one or more microprocessors, memory, storage, and network controllers and runs its own operating system and applications.

concurrent copy

A function of the DFSMSdss component that is used to back up any collection of data at a point in time with minimum down time for the database or application that uses the collection of data.

consistency group

A group of copy relationships between virtual volumes or data sets that are maintained with the same time reference so that all copies are consistent in time.

copyback

A process that moves data back to its expected or preferred location to maintain an array in a more efficient configuration after a failed drive is replaced.

copy set

The set of source volumes or target volumes involved in a FlashCopy® operation.

count key data

1. An architecture for a direct access storage device (DASD) device or logical device that specifies the access mechanisms for the logical data units on the device through a specific set of supported channel commands. Extensions to the CKD command set form the basis of Extended CKD.
2. A data recording format that uses self-defining record formats in which each record on a volume is represented by up to three fields: a count field identifying the record and specifying its format, an optional key field that can be used to identify the data area contents, and an optional data field that typically contains the user data. See also [data record](#), [storage architecture type](#).

count-key-data record (CKD record)

See [data record](#).

CRU

See [customer-replaceable unit](#).

customer-replaceable unit (CRU)

An assembly or part that can be replaced in its entirety by a user when any one of its components fails.

cylinder

A unit of storage on a count-key-data (CKD) device with a fixed number of tracks.

D

data collection

The process of obtaining performance and availability monitoring data and providing that data to a metric evaluator. Examples of data collectors include Domain Name System (DNS) probes, web page analyzers, or database analyzers. See also [discovery](#).

data consistency

A characteristic of the data at the target site where dependent write order is maintained to guarantee the recoverability of applications.

data record

A basic unit of data recording format. See also [count key data](#), [fixed-block architecture](#).

data reduction

A set of techniques that can be used to reduce the amount of usable capacity that is required to store data. Examples of data reduction include data deduplication and compression. See also [data reduction savings](#), [stored capacity](#).

data reduction savings

The total amount of usable capacity that is saved in a system, pool, or volume through the application of an algorithm such as compression or deduplication on the written data. This saved capacity is the difference between the written capacity and the used capacity. See also [data reduction](#).

data source

A storage resource or agent that provides data about a storage environment.

destage

To move data from cache to a nonvolatile storage medium.

discovery

The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved. See also [data collection](#).

distributed RAID

An alternative RAID scheme where the number of drives that are used to store the array can be greater than the equivalent, typical RAID scheme. The same data stripes are distributed across a greater number of drives, which increases the opportunity for parallel I/O and hence improves overall array performance. See also [rebuild area](#).

DNS

See [Domain Name System](#).

Domain Name System (DNS)

The distributed database system that maps domain names to IP addresses.

drive

A data storage device. A drive can be either a magnetic disk drive or a solid-state drive (SSD).

drive class

A combination of drive technology and speed, which uniquely defines a class of drives that have approximately the same performance characteristics.

drive technology

A category of a drive that pertains to the method and reliability of the data storage techniques being used on the drive. Possible values include enterprise (ENT) drive, nearline (NL) drive, or solid-state drive (SSD).

E

ECKD

See [extended count key data](#).

effective capacity

The amount of provisioned capacity that can be created in a system or pool without running out of usable capacity given the current data reduction savings being achieved. This capacity equals the usable capacity divided by the data reduction savings percentage.

enclosure

The metal structure in which various electronic components are mounted.

encryption deadlock

The inability to access encryption keys to decrypt data. See also [encryption recovery key](#).

encryption key label

The list of encryption key labels used by the storage system to identify keys that will be used on the key server.

encryption key manager

See [encryption key server](#).

encryption key server

An internal or external system that runs a key manager that receives and then serves existing encryption keys or certificates to a storage system.

encryption recovery key

An encryption key that allows a method to recover from an encryption deadlock situation where the normal encryption key servers are not available. See also [encryption deadlock](#).

enterprise

Pertaining to a type of data storage device that has higher error recovery limits, vibration tolerance, and end-to-end error detection than standard desktop hard drives.

enterprise repository

A component of the data server that records and stores all information about the monitored computers' storage assets and their usage over time. The repository is organized into relational database tables and is accessed by the data server using Java™ Database Connectivity (JDBC).

event

An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process.

extended count key data (ECKD)

An extension of the count-key-data (CKD) architecture. It includes additional commands that can be used to improve performance.

extent type

See [storage architecture type](#).

F

failback

The restoration of an appliance to its initial configuration after detection and repair of a failed network or component.

failover

An automatic operation that switches to a redundant or standby system or node in the event of a software, hardware, or network interruption.

FB

See [fixed block](#).

FBA

See [fixed-block architecture](#).

FC

See [Fibre Channel](#).

FC-AL

See [Fibre Channel Arbitrated Loop](#).

FCIP

See [Fibre Channel over IP](#).

FCP

See [Fibre Channel Protocol](#).

feature activation code

See [license key](#).

Fibre Channel (FC)

A technology for transmitting data between computer devices. It is especially suited for attaching computer servers to shared storage devices and for interconnecting storage controllers and drives. See also [zoning](#).

Fibre Channel Arbitrated Loop (FC-AL)

An implementation of the Fibre Channel standards that uses a ring topology for the communication fabric; refer to American National Standards Institute (ANSI) INCITS 272-1996, (R2001). In this topology, two or more Fibre Channel end points are interconnected through a looped interface.

Fibre Channel connection (FICON®)

A Fibre Channel communication protocol designed for IBM mainframe computers and peripherals.

Fibre Channel extender

A device used to extend a Fibre Channel link over a greater distance than is supported by the standard, usually a number of miles or kilometers. Devices must be deployed in pairs at each end of a link.

Fibre Channel over IP (FCIP)

A network storage technology that combines the features of the Fibre Channel Protocol and the Internet Protocol (IP) to connect distributed SANs over large distances.

Fibre Channel Protocol (FCP)

The serial SCSI command protocol used on Fibre Channel networks. See also [open system](#).

FICON

See [Fibre Channel connection](#).

field-replaceable unit (FRU)

An assembly that is replaced in its entirety when any one of its components fails.

file module

A component that provides file systems to network users. A file module must be provided with storage for the file systems.

fileset

See [file set](#).

file set

1. A subset of a file system that provides granularity of functions such as snapshots or quotas within the file system.
2. A hierarchical grouping of files managed as a unit for balancing workload across a cluster.

file system (FS)

A collection of files and certain attributes associated with those files.

file system storage

Data storage that is organized into files and directories.

fixed block (FB)

See [fixed-block architecture](#).

fixed-block architecture (FBA)

An architecture for a virtual device that specifies the format of and access mechanisms for the virtual data units on the device. The virtual data unit is a block. All blocks on the device are the same size (fixed size). The system can access them independently. See also [data record](#), [storage architecture type](#).

FlashCopy

1. Pertaining to a point-in-time copy where a virtual copy of a volume is created. The target volume maintains the contents of the volume at the point in time when the copy was established. Any subsequent write operations to the source volume are not reflected on the target volume.
2. An optional feature of the Storage System DS family that can make an instant copy of data, that is, a point-in-time copy of a volume.

flash drive

A data storage device, which is typically removable and rewritable, that uses solid-state memory to store persistent data. See also [flash module](#).

flash module

A modular hardware unit containing flash memory, one or more flash controllers, and associated electronics. See also [flash drive](#).

flush-through mode

See [write-through mode](#).

form factor

The industry-standard physical dimensions of a storage system drive enclosure. Possible values include “3.5 inch”, “2.5 inch”, and “1.8 inch.”

frame

The hardware support structure, covers, and all electrical parts mounted therein that are packaged as one entity for shipping.

freeze

An operation in which a storage system blocks I/O from the host system to the affected volumes on the primary site. A freeze operation stops mirroring between the primary and secondary volumes to ensure data consistency at the secondary site. See also [thaw](#).

FRU

See [field-replaceable unit](#).

FS

See [file system](#).

full restore operation

A copy operation where a local volume is created by reading an entire a volume snapshot from cloud storage.

full snapshot

A type of volume snapshot that contains all the volume data. When a full snapshot is created, an entire copy of the volume data is transmitted to the cloud.

G

General Parallel File System (GPFS)

A high-performance shared-disk file system that can provide data access from nodes in a clustered system environment.

Global Copy

A non-synchronous long-distance copy option for data migration and backup. See also [remote mirror and copy](#).

globally unique identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

Global Mirror

A method of an asynchronous replication that maintains data consistency across multiple volumes within or across multiple systems. Global Mirror is generally used where distances between the source site and target site cause increased latency beyond what the application can accept.

GPFS

See [General Parallel File System](#).

grain size

The unit size for allocating space on thin-provisioned volumes, such as 32, 64, 128, and 256 kibibyte. The grain size is defined when a volume is created.

GUID

See [globally unique identifier](#).

H

Hardware Management Console (HMC)

1. In a system storage environment, a system that acts as the focal point for configuration, management of Copy Services functions, and maintenance.
2. A system that controls managed systems, including the management of logical partitions and use of Capacity Upgrade on Demand. Using service applications, the HMC communicates with managed systems to detect and consolidate information, which can then be sent for analysis.

heat map

An overlay that shows where the user performed actions on a web page. The data is summarized and shown as a colored marker at the location of the user action over a snapshot of the web page. The overlay lets the user see where and how the web page user interacted with the web page.

HMC

See [Hardware Management Console](#).

host

A physical or virtual computer system that hosts computer applications, with the host and the applications using storage. See also [client](#), [host](#), [server](#).

host cluster

A configured set of physical or virtual hosts that share one or more storage volumes in order to increase scalability or availability of computer applications.

host interface card

See [interface card](#).

host object

A logical representation of a host within a storage system that is used to represent the host for configuration tasks.

host volume

A volume that represents the volume functional role from an application point of view. The host volume can be connected to a host or server. It receives read, write, and update application I/O, depending on the site to which the application is writing.

hot-spare

Pertaining to redundant hardware (such as an adapter, a disk, a drive, or a server) that is installed and available in the event of a hardware failure.

HyperSwap®

Pertaining to a function that provides continuous, transparent availability against storage errors and site failures, and is based on synchronous replication.

hypervisor

Software or a physical device that enables multiple instances of operating systems to run simultaneously on the same hardware.

I

incremental restore operation

A copy operation where a local volume is modified to match a volume snapshot by reading from cloud storage only the parts of the volume snapshot that differ from the local volume.

incremental snapshot

A type of volume snapshot where the changes to a local volume relative to the volume's previous snapshot are stored on cloud storage.

input/output (I/O)

Pertaining to a device, process, channel, or communication path involved in data input, data output, or both.

interface card

An optional part of a node canister that provides the system with additional host and storage connectivity options.

interface node

A node that connects a system to an Internet Protocol (IP) network for file-serving capabilities by using service protocols.

Internet Small Computer System Interface (iSCSI)

An IP-based standard for linking data storage devices over a network and transferring data by carrying SCSI commands over IP networks. See also [Small Computer System Interface](#).

I/O

See [input/output](#).

I/O enclosure

A hardware unit in a storage system where data is transferred into and out of the system.

iSCSI

See [Internet Small Computer System Interface](#).

J

journal volume

A volume that holds a consistent copy of data until a new consistent copy is formed. The journal volume restores the last consistent point during a recovery.

K

key server

1. A server that negotiates the values that determine the characteristics of a dynamic virtual private network (VPN) connection that is established between two endpoints.
2. See [encryption key server](#).

L

launch-in-context

An operation in which a user starts a secondary application from a primary application to perform a specific task. Using the parameters, navigation instructions, and user credentials that are supplied by the primary application, the secondary application opens to the specific place in which to complete the task.

licensed capacity

The amount of capacity on a storage system that a user is entitled to configure.

license key

An alphanumeric code that activates a licensed function on a product.

license key file

A file that contains one or more licensed keys.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

LUN

See [logical unit number](#).

M

machine signature

A string of characters that identifies a system. A machine signature might be required to obtain a license key.

Management Information Base (MIB)

In the Simple Network Management Protocol (SNMP), a database of objects that can be queried or set by a network management system.

management node

A node that is used for configuring, administering, and monitoring a system.

management server

A system that provides a central point of control for managing data replication.

management server relationship

A connection between two replication servers, where one server acts as the active server and replicates the data that is necessary for the standby server to take control of the replication environment.

maximum replication delay

The number of seconds that Metro Mirror or Global Mirror replication can delay a write operation to a volume.

Metro Global Mirror

1. A cascaded solution where Metro Mirror synchronously copies data to the target site. This Metro Mirror target is the source volume for Global Mirror that asynchronously copies data to a third site. This solution has the potential to provide a disaster recovery with no data loss at Global Mirror distances when the intermediate site does not participate in the disaster that occurs at the production site.
2. A three-site, high availability, disaster recovery solution. Metro Global Mirror uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

Metro Mirror

A method of synchronous replication that maintains data consistency across multiple volumes within the system. Metro Mirror is generally used when the write latency caused by the distance between the source site and target site is acceptable to application performance.

MIB

See [Management Information Base](#).

N

namespace

The scope within which a Common Information Model (CIM) schema applies.

native interface

An interface that is specific to a system or subsystem.

nearline

Pertaining to a type of storage in which data is available in a short amount of time, but not instantly.

nearline SAS drive

A drive that combines the high capacity data storage technology of a Serial Advanced Technology Attachment (SATA) drive with the benefits of a serial-attached SCSI (SAS) interface for improved connectivity.

node

A single processing unit within a system. For redundancy, multiple nodes are typically deployed to make up a system.

O

open system

A system that complies with industry-defined interoperability standards. An open system can be connected to other systems complying with the same standards.

See also [Fibre Channel Protocol](#), [Small Computer System Interface](#).

order confirmation code

See [authorization code](#).

overhead capacity

An amount of usable capacity that is occupied by metadata in a system or pool and other data that is used for system operations.

overprovisioned ratio

The ratio of provisioned capacity to usable capacity in a system or pool.

overprovisioning

The result of creating more provisioned capacity in a storage system or pool than there is usable capacity. Overprovisioning occurs when thin provisioning or data reduction techniques ensure that the used capacity of the provisioned volumes is less than their provisioned capacity.

P

parent pool

A storage pool that receives its capacity from MDisk and has, or will have, some of its capacity allocated to child pools. See also [child pool](#).

performance group

A collection of volumes that is assigned the same performance characteristics. See also [performance policy](#).

performance policy

A policy that specifies performance characteristics, for example quality of service (QoS). See also [performance group](#).

PFC

See [priority flow control](#).

ping

1. The command that sends an Internet Control Message Protocol (ICMP) echo-request packet to a gateway, router, or host with the expectation of receiving a reply.
2. A job that tracks the availability of assets and that is performed by an agent. Several ping jobs can be used to monitor the availability of any computer or subset of computers in the network.

pool

1. See [storage pool](#).
2. A grouping of storage space that consists of volumes, logical unit numbers (LUNs), or addresses that share a common set of administrative characteristics.

pool pair

Two storage pools that are required to balance workload. Each storage pool is controlled by a separate node.

port

The physical entity within a host, system, or storage system that performs the data communication (transmitting and receiving) over the Fibre Channel.

practice volume

A volume that can be used to test disaster-recovery actions while maintaining disaster-recovery capability.

primary site

A physical or virtual site that is made up of hardware, network, and storage resources. Typically, production operations run at the primary site. Data can be replicated to a secondary site for disaster recovery and failover operations. See also [secondary site](#).

primordial pool

Storage capacity that is unallocated on a storage device. Storage pools are created by allocating storage capacity from primordial pools.

priority flow control (PFC)

A link-level flow control mechanism, IEEE standard 802.1Qbb. PFC operates on individual priorities. Instead of pausing all traffic on a link, PFC is used to selectively pause traffic according to its class.

probe

A data collection job that itemizes and creates an inventory of assets, such as computers, controllers, disk drives, file systems, and logical units.

projected capacity

The estimated volume capacity that is available for volume creation, given the current average performance of any data compression, excluding thin-provisioning savings. See also [thin-provisioning savings](#).

protocol

A set of rules controlling the communication and transfer of data between two or more devices or systems in a communication network.

provisioned capacity

The total capacity of all volumes and volume copies in a system or pool.

R

- rack
 - A free-standing structure that can hold multiple servers, storage systems, chassis, switches, and other devices.
- RAID
 - See [Redundant Array of Independent Disks](#).
- RAID 0
 - A data striping technique, which is commonly called RAID Level 0 or RAID 0 because of its similarity to common, RAID, data-mapping techniques. It includes no data protection, however, so, strictly speaking, the appellation RAID is a misnomer. RAID 0 is also known as data striping.
- RAID 1
 - A form of storage array in which two or more identical copies of data are maintained on separate media.
- RAID 10
 - A collection of two or more physical drives that present to the host an image of one or more drives. In the event of a physical device failure, the data can be read or regenerated from the other drives in the RAID due to data redundancy.
- RAID 5
 - A form of parity RAID in which the disks operate independently, the data stripe size is no smaller than the exported block size, and parity check data is distributed across the array's disks.
- RAID 6
 - A form of RAID that can continue to process read and write requests to all of an array's virtual disks in the presence of two concurrent disk failures.
- RAID level
 - The level of protection provided by the specific techniques of striping, mirroring, or parity used by a Redundant Array of Independent Disks (RAID).
- RAID type
 - See [RAID level](#).
- raw capacity
 - The reported capacity of the drives in the system before formatting or RAID (Redundant Array of Independent Disks) is applied.
- rebuild area
 - Reserved capacity that is distributed across all drives in a redundant array of drives. If a drive in the array fails, the lost array data is systematically restored into the reserved capacity, returning redundancy to the array. The duration of the restoration process is minimized because all drive members simultaneously participate in restoring the data. See also [distributed RAID](#).
- reclaimable capacity
 - The amount of provisioned capacity that can be recovered without causing stress or performance degradation. This capacity identifies the amount of resources that can be reclaimed and provisioned to other objects in an environment.
- reclaimed capacity
 - See [reclaimable capacity](#).
- recovery key
 - See [encryption recovery key](#).
- recovery point objective
 - The maximum amount of data loss that can be tolerated during a service interruption.
- Redundant Array of Independent Disks (RAID)
 - A collection of two or more physical disk drives that present to the host an image of one or more logical disk drives. In the event of a physical device failure, the data can be read or regenerated from the other disk drives in the array due to data redundancy.
- remote mirror and copy
 - A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also [Global Copy](#).
- repo
 - See [repository](#).
- repository (repo)
 - A persistent storage area for data and other application resources.
- reserved capacity
 - The amount of used capacity that is made up of capacity reserved for system use. See also [total usable physical capacity](#).
- resource
 - In a storage environment, an entity that is monitored. Resources can include fabrics, switches, computers, and storage systems.
- role
 - A job function that identifies the tasks that a user can perform and the resources to which a user has access. A user can be assigned one or more roles.
- role pair
 - The association of two volume roles in a session that take part in a copy relationship. For example, in a Metro Mirror session, the role pair can be the association between host volumes at the primary site and host volumes at the secondary site (H1-H2).

S

- scan
 - A data collection job that monitors storage usage and file statistics on the resources in an environment.
- SCSI
 - See [Small Computer System Interface](#).
- SCSI device
 - A product, such as a drive or adapter, connected to a host through an I/O interface using the Small Computer System Interface (SCSI) protocol. A SCSI device is either an initiator, target, or both. See also [Small Computer System Interface](#).
- SCSI-FCP
 - See [SCSI Fibre Channel Protocol](#).
- SCSI Fibre Channel Protocol (SCSI-FCP)
 - A standard that defines the protocol used to transfer Small Computer System Interface (SCSI) commands over the transport physical layer of the Fibre-Channel interface. This standard is published by ANSI as X3.269-1996.
- SCSI initiator
 - The system component that initiates communications with attached targets.
- SCSI target
 -

A device that acts as a subordinate to a SCSI initiator and consists of a set of one or more logical units (LUs), each with an assigned logical unit number (LUN). The LUs on the SCSI target are typically I/O devices.

secondary site
A physical or virtual site that is made up of the hardware, network, and storage resources that support the recovery needs of the primary site. When a failure occurs at the primary site, operations can continue at the secondary site. See also [primary site](#).

server
A computer program or a device that provides functions for other programs or devices, called clients. See also [client](#), [host](#).

Server Message Block (SMB)
A protocol that manages requests and responses in a client/server environment so that clients on a network can share files, directories, and devices. See also [Server Message Block 2.0](#).

Server Message Block 2.0
A higher performing, more scalable version of Server Message Block (SMB). This protocol can send multiple commands in the same packet and uses larger buffer sizes. See also [Server Message Block](#).

session
A collection of source and target volumes that are managed to create consistent copies of data. The type of data replication that is associated with the session determines the actions that can be conducted for the volumes.

site awareness
The association of a location with each storage system in a session. Site awareness ensures that a volume can only be selected for a session if it matches the location of the site. Site awareness helps to prevent both reversing a hardware relationship and selecting volumes at the wrong location.

Small Computer System Interface (SCSI)
An ANSI-standard electronic interface that allows personal computers to communicate with peripheral hardware, such as disk drives, tape drives, CD-ROM drives, printers, and scanners faster and more flexibly than previous interfaces. See also [Internet Small Computer System Interface](#), [open system](#), [SCSI device](#).

SMB
See [Server Message Block](#).

SMI-S
See [Storage Management Initiative Specification](#).

solid-state drive (SSD)

1. A storage device that contains nonvolatile flash memory. A solid-state drive (SSD) has no moving mechanical components.
2. See [flash drive](#).

space
See [capacity](#).

space efficient
See [thin provisioning](#).

spare drive
A drive reserved in an array for rebuilding a failed drive in a RAID. Should a drive fail in a RAID, a spare drive from within that device adapter (DA) pair will be selected to rebuild it.

SSD
See [solid-state drive](#).

standard-provisioned volume
A volume that completely uses storage at creation.

standard provisioning
The ability to completely use a volume's capacity for that specific volume.

standby management server
A management server that is a backup for the active server. The replication environment cannot be monitored or managed from the standby server.

storage architecture type (storage type)
The type of storage architecture, either count key data (CKD) or fixed block (FB), for which an array, pool, or volume is provisioned. See also [count key data](#), [fixed-block architecture](#).

storage enclosure
A specialized chassis that is designed to hold and power drives while providing a mechanism to allow them to communicate to one or more separate computers.

Storage Management Initiative Specification (SMI-S)
A design specification developed by the Storage Networking Industry Association (SNIA) that specifies a secure and reliable interface with which storage management systems (SMSs) can identify, classify, monitor, and control physical and logical resources in a storage area network (SAN). The interface integrates the various devices to be managed in a SAN and the tools used to manage them.

storage node
A component of a storage system that provides internal storage or a connection to one or more external storage systems.

storage pod

1. A subcomponent of a network-attached storage (NAS) system that consists of two or more storage nodes and one or more supported storage systems.
2. A logical entity of components of a system consisting of two storage nodes, and one or two storage subsystems directly connected with these storage nodes.

storage pool (pool)
A collection of storage that identifies an underlying set of resources. These resources provide the capacity and management requirements for a volume or set of volumes.

storage resource group
A named collection of logically related resources that are monitored by Tivoli® Storage Productivity Center. Monitored resources can include fabrics, switches, computers, storage systems, and other storage resource groups.

storage system
A system that provides persistent storage within a network. A storage system can include facilities for host attachment, user role authentication, a command-line interface (CLI), a graphical user interface (GUI), and storage devices that most often include Redundant Array of Independent Disks (RAID) controllers. It might also include agents for enabling third-party management software to monitor or manage the storage devices.

storage type
See [storage architecture type](#).

stored capacity
The amount of capacity that is used to store data that is written by a host after data reduction. See also [data reduction](#), [total usable physical capacity](#).

support assistance
A function that is used to provide support personnel access to the system to complete troubleshooting and maintenance tasks.

synchronous replication
A type of replication in which the application write operation is made to both the source volume and target volume before control is given back to the application. See also [asynchronous replication](#).

syslog

A standard for transmitting and storing log messages from many sources to a centralized location to enhance system management.

T

target volume

A volume that receives data from a host volume or another intermediate volume.

thaw

An operation in which a storage system releases the block of the I/O from the host system to the affected volumes on the primary site. A thaw operation can occur after a freeze operation ends and consistency is formed on the secondary site. See also [freeze](#).

thin-provisioned volume

A volume that allocates storage when data is written to it.

thin provisioning

1. The ability to defer capacity allocation on a storage resource until data is actually written to it.
2. A mechanism that provides the ability to define logical volume sizes that are larger than the physical capacity installed on the system.

thin-provisioning savings

The total amount of usable capacity that is saved in a system, pool, or volume by consuming usable capacity only when needed as a result of write operations. The capacity that is saved is the difference between the provisioned capacity minus the written capacity. See also [projected capacity](#), [volume capacity](#), [written capacity](#).

total capacity savings

The total amount of usable capacity that is saved in a system, pool, or volume through thin provisioning and data reduction techniques. This saved capacity is the difference between the used usable capacity and the provisioned capacity.

total usable physical capacity

The amount of physical configured storage space that is available for stored capacity or reserved capacity. This capacity can consist of both internal storage through arrays and external storage through MDisk. See also [capacity threshold](#), [reserved capacity](#), [stored capacity](#).

track space-efficient volume (TSE volume)

A volume in which storage space is allocated on an as-needed basis by using space on the target volume only when tracks are copied from the source volume to the target volume.

transparent cloud tiering

The functions that use cloud storage as an extension of on-premises storage.

trial license

A temporary entitlement to use a licensed function.

TSE for FlashCopy

A thin-provisioning method in which storage space is allocated from a TSE repository on an as needed basis. See also [TSE repository](#).

TSE repository

The amount of capacity in a storage pool reserved for volumes that use a thin-provisioning method of TSE for FlashCopy. See also [TSE for FlashCopy](#).

TSE volume

See [track space-efficient volume](#).

U

unmapped volume capacity

The amount of volume capacity that is not mapped to a host. See also [volume capacity](#).

update

1. To apply fixes to a system.
2. Software maintenance such as a manufacturing refresh, refresh pack, or fix pack that changes the modification level of a product.
3. To modify a file or data set with current information.

upgrade

1. Any hardware or software change to a later release, or any hardware addition or software addition.
2. To install a new version or release of a product to replace an earlier version or release of the same product.

usable capacity

The amount of capacity that is provided for storing data on a system, pool, array, or MDisk after formatting and RAID techniques are applied.

used capacity

The amount of usable capacity that is taken up by data or overhead capacity in a system, pool, array, or MDisk after data reduction techniques have been applied.

user role

An identifier that is assigned to a user that defines the set of permissions that are granted to that user.

V

virtual capacity

See [provisioned capacity](#).

virtualization

The substitution of virtual resources for actual resources, where the virtual resources have the same functions and external interfaces as their counterparts, but differ in attributes, such as size, performance, and cost. Virtualization is commonly applied to physical hardware resources by combining multiple physical resources into shared pools from which users receive virtual resources.

virtualized capacity

The amount of capacity that is contributed to a storage pool by a given provisioning group.

virtual machine (VM)

An emulation of a particular computer system. Virtual machines operate based on the computer architecture and functions of a real or hypothetical computer. Their implementations might involve specialized hardware, software, or a combination of both.

virtual storage area network (VSAN)

A fabric within the storage area network (SAN).

VM

See [virtual machine](#).

volume

A fixed amount of physical or virtual storage on a data storage medium.

volume access set

The set of I/O groups that allows host access to a volume. This set can optionally include the caching I/O group.

volume capacity

The total capacity for all volumes in a system or storage pool. Volume capacity is defined by the client when a volume is created and surfaced to the host. See also [thin-provisioning savings](#), [unmapped volume capacity](#).

volume snapshot

A collection of objects on a cloud storage account that represents the data of a volume at a particular time.

VSAN

See [virtual storage area network](#).

W

worldwide ID (WWID)

A name identifier that is unique worldwide and that is represented by a 64-bit value that includes the IEEE-assigned organizationally unique identifier (OUI).

worldwide name (WWN)

A 64-bit, unsigned name identifier that is unique.

worldwide node name (WWNN)

A unique 64-bit identifier for a host containing a Fibre Channel port. See also [worldwide port name](#).

worldwide port name (WWPN)

A unique 64-bit identifier associated with a Fibre Channel adapter port. The WWPN is assigned in an implementation-independent and protocol-independent manner. See also [worldwide node name](#).

write-through mode

A process in which data is written to a storage device at the same time as the data is cached.

written capacity

The amount of usable capacity that would have been used to store written data in a system or pool if data reduction was not applied. See also [thin-provisioning savings](#).

written capacity limit

The largest amount of capacity that can be written to a drive, array, or MDisk. The limit can be reached even when usable capacity is still available.

WWID

See [worldwide ID](#).

WWN

See [worldwide name](#).

WWNN

See [worldwide node name](#).

WWPN

See [worldwide port name](#).

Z

z Global Mirror

A method of an asynchronous replication function that maintains data consistency across multiple volumes that are attached to a z/OS® system. Time-based data consistency is maintained through the Data Facility Storage Management Subsystem (DFSMS) system data mover (SDM) component.

zone

A logical grouping of switches, switch ports, and their attached devices in a fabric.

zone alias

A name that is given to a collection of one or more zone members to be managed together.

zone set

A group of zones that function together on a fabric.

zoning

The grouping of multiple ports to form a virtual, private, storage network. Ports that are members of a zone can communicate with each other, but are isolated from ports in other zones. See also [Fibre Channel](#).